INNOVATIONS IN AUDIO • AUDIO ELECTRONICS • THE BEST IN DIY AUDIO

www.audioxpress.com

audioxpress

ADVANCING THE EVOLUTION OF AUDIO TECHNOLOGY

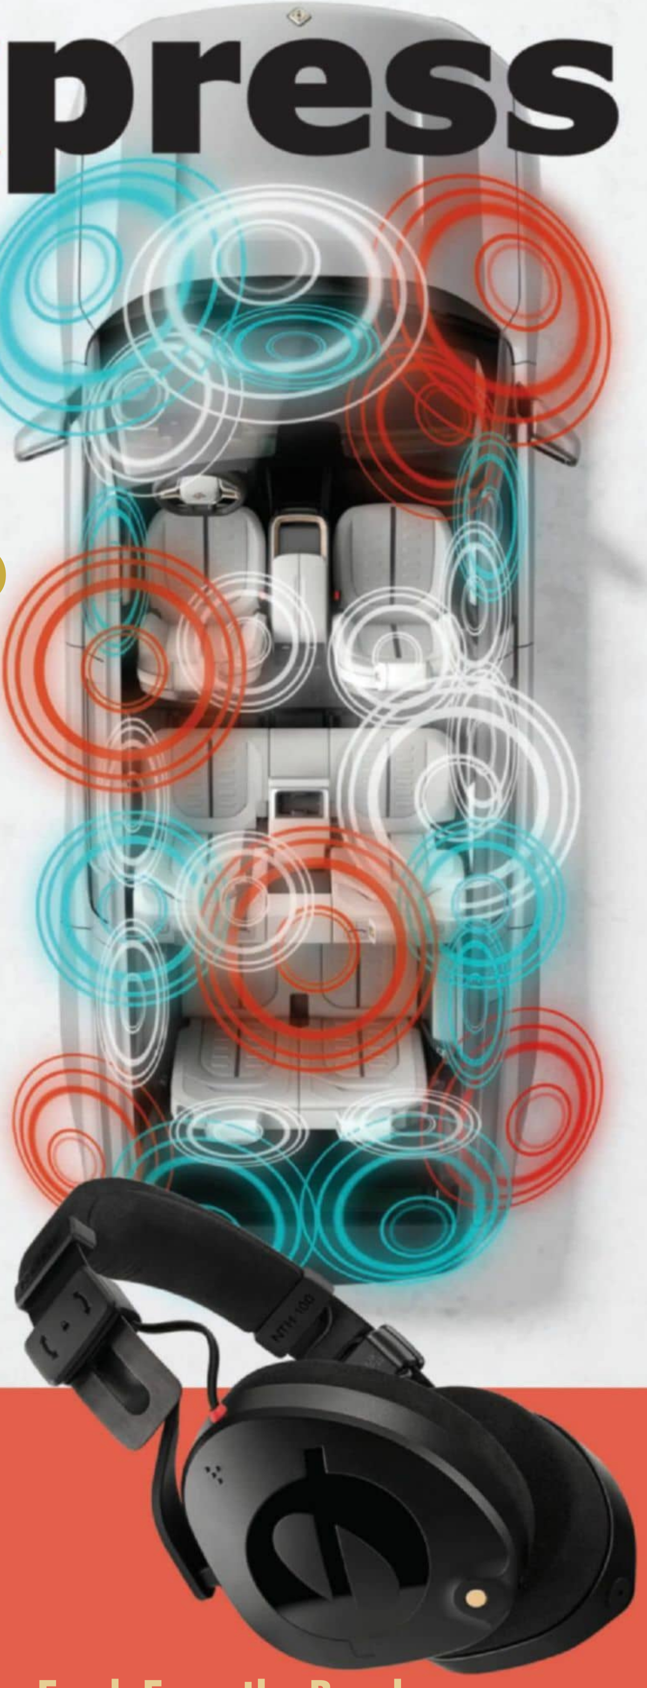

Automotive Audio

Audio for All-Electrical Vehicles
By J. Martins

Upmixing and the Future of In-Car Audio
By E. Hemström and R. Fleischer (Dirac)

Evaluation of Automotive Audio Energy Efficiency
By Philipp Paul Klose

The John Whitecar Interview
By Adam Levenson

Practical Test & Measurement

ANC Headset Testing
By H. W. Gierlich (HEAD acoustics)

Development and Design

Signal Processing and Filter Bank Design for Hearing Aids
By Dana Helmink (Widex)

You Can DIY!

Building a Guitar-Controlled Synthesizer: LFO and ADSR
By Ethan Winer

Sound Control

The Acoustical Effects of Room Shapes
By Richard Honeycutt

Fresh From the Bench

Røde NTH-100 Headphones
By Brent Butterworth

$9.00US $10.00CAN

JUNE 2022

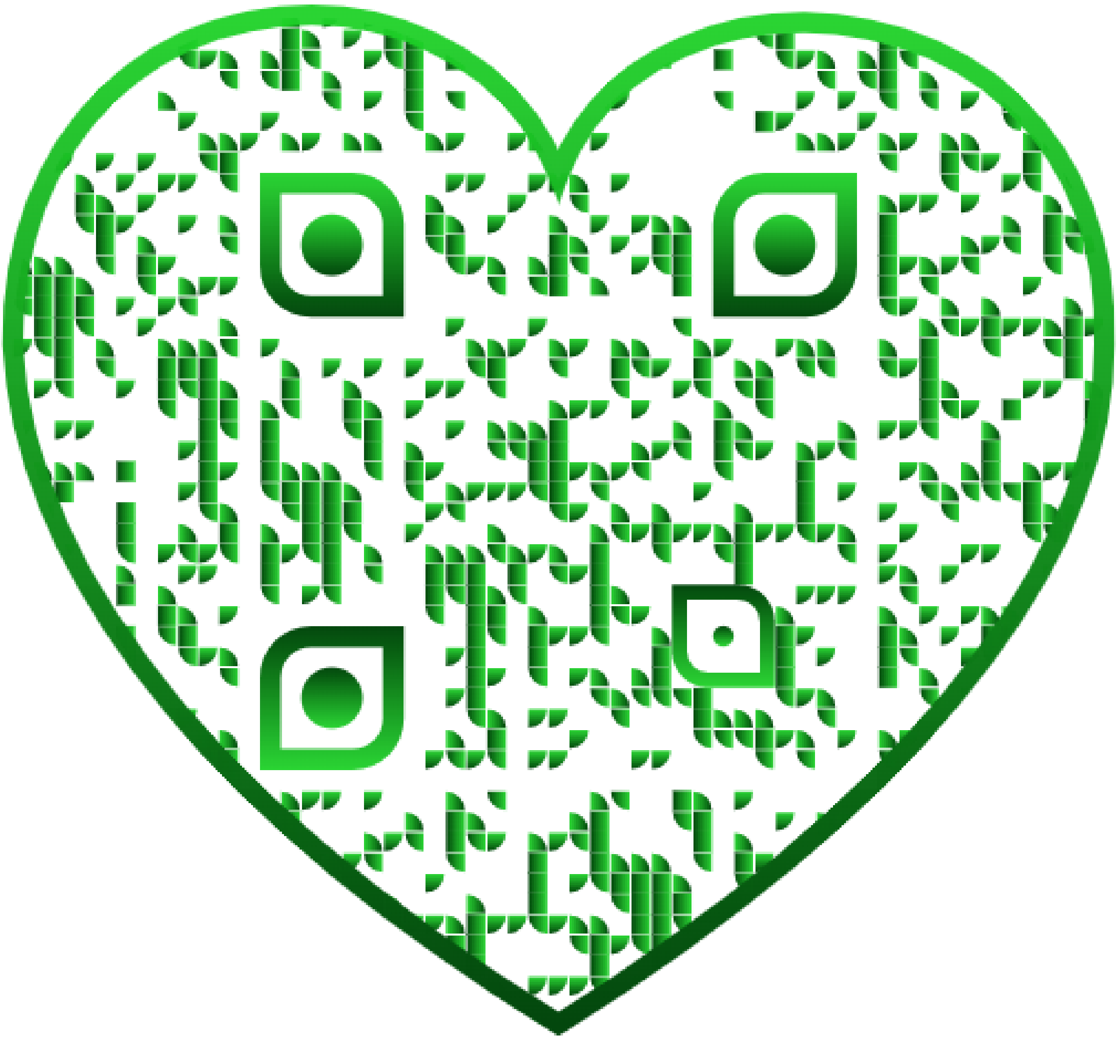

ANATOMY OF A 300B

PRESENTED BY *Western Electric*

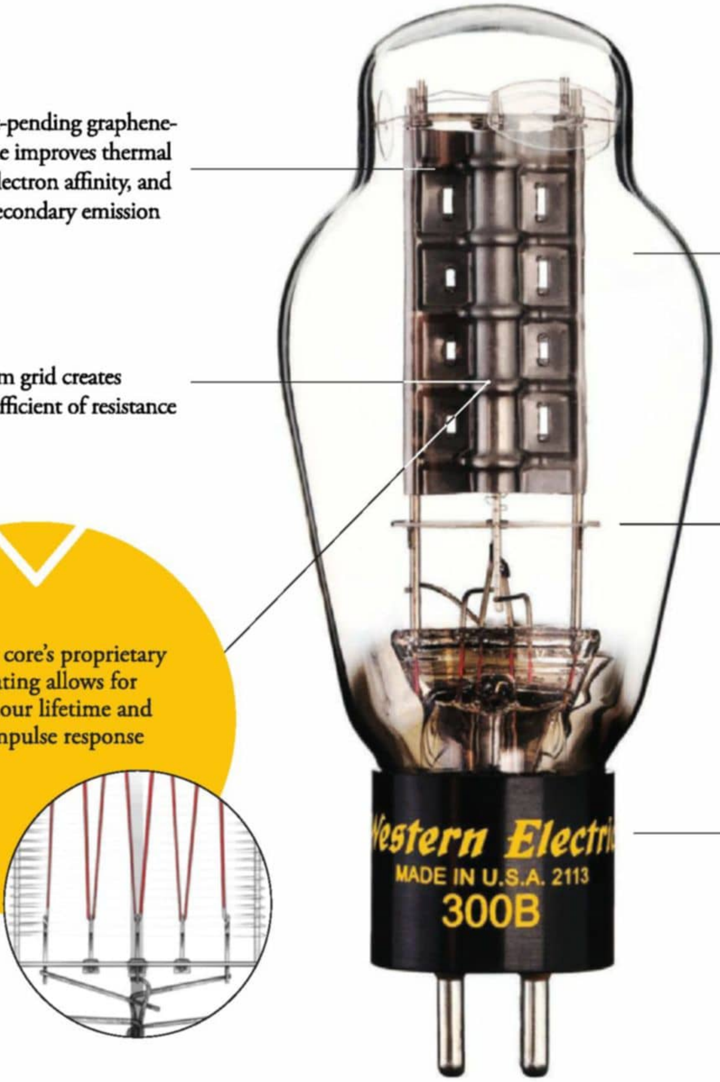

New, patent-pending graphene-coated anode improves thermal emissivity, electron affinity, and eliminates secondary emission

Molybdenum grid creates negative coefficient of resistance

Filament core's proprietary oxide coating allows for 40,000 hour lifetime and superb impulse response

Ultra-high vacuum finishing at 2×10^{-7} Torr, made possible with state-of-the-art turbomolecular pumping systems

Proprietary mica coating increases dielectric strength, optimizes higher voltage operation

Quality guaranteed with a five-year extended warranty

There's more to it than cathode, anode, and grid. Each Western Electric 300B is assembled at our new facility by the hands of our highly trained staff using the highest quality materials and modern technology—all while honoring the original 1938 specifications and tooling.

What's the result? What's in it for the listener? Unmatched manufacturing quality and the most powerful, full-bodied sound there is. When you're through compromising, order a matched pair online or from Authorized Dealers worldwide.

AVAILABLE NOW

westernelectric.com/300b

WESTERN ELECTRIC • MAKER OF ELECTRON TUBES & HIGH FIDELITY • 201 W GORDON AVE • ROSSVILLE, GA • USA • +1 404.352.2000

June 2022 ISSN 1548-0628

www.audioxpress.com

audioXpress (US ISSN 1548-0628) is published monthly, at $50 per year for the US, at $65 per year for Canada, and at $75 per year Foreign/ROW, by KCK Media Corp., P.O. Box 417, Chase City, VA 23924, US Periodical Postage paid at Chase City, VA, and additional offices.

Head Office:
KCK Media Corp.
P.O. Box 417
Chase City, VA 23924, US
Phone: 434-533-0246
Fax: 888-980-1303

Subscription Management:
audioXpress
P.O. Box 417
Chase City, VA 23924, US
Phone: 434-533-0246
E-mail: custserv@audioxpress.com
Internet: www.audioxpress.com

Postmaster: Send address changes to:
audioXpress
P.O. Box 417
Chase City, VA 23924, US

Advertising:
Strategic Media Marketing, Inc.
2 Main Street
Gloucester, MA 01930, US
Phone: 978-281-7708
Fax: 978-281-7706
E-mail: audioxpress@smmarketing.us
Advertising rates and terms available on request.

Editorial Inquiries:
Send editorial correspondence, suggestions, and contributed articles to:
audioXpress, Editorial Department
P.O. Box 417
Chase City, VA 23924, US
E-mail: editor@audioxpress.com

Legal Notice:
Each design published in audioXpress is the intellectual property of its author and is offered to readers for their personal use only.
Any commercial use of such ideas or designs without prior written permission is an infringement of the copyright protection of the work of each author.

© KCK Media Corp. 2022
Printed in the US

Driving Audio Innovations

This issue of *audioXpress* focuses on the vibrant segment of automotive audio, which is where a lot of exciting technologies are being explored. Automotive cabin acoustics are not an easy environment but when we combine the latest transducers, amplifiers, and DSP, with careful design and tuning by system experts that now have access to the most sophisticated computer simulation and measurement tools, it's easy to understand why cars are currently at the cutting edge of applications in audio.

The transition to electric engines has completely changed the possibilities for car sound, as our annual Market Update details. We are now able to have consistent sound experiences inside the latest cars that the vast majority of people will never be able to experience elsewhere. And that's due to the fact that audio has now become a top priority inside electric vehicles (EVs). No longer as a mere extension of infotainment but as the essential component for the whole experience inside the vehicle.

Automotive audio systems in EVs need to cancel undesirable road noise; create suitable acoustic environments for passenger activities, including managing separate zones within the cabin; provide the most effective communications; offer the best possible musical listening experience for all passengers; and support voice control at all moments. Sound systems inside the car are at the core of all user interactions with the vehicle, including creating an audible human-interface for the car motion and propulsion engine (given that EVs are essentially silent), while enabling users to interact with the car systems through voice. And outside the car, a new class of audio systems is required to alert pedestrians to the presence of silent EVs.

Along with the smart home, this integrated automotive approach is one of the most formidable opportunities ever for the audio industry. It allows new companies—particularly speaker companies—a chance to prove concepts that would otherwise struggle to demonstrate their worth, and at a scale that would otherwise make investments too risky.

It also enables software companies, such as DSP Concepts, Meridian, Dirac, and even Dolby Laboratories, to develop technologies that are now at the core of automotive systems. And very likely, those investments will translate to benefits for many other consumer and even professional audio applications.

The fact that automotive systems now use multiple microphones (and other sensors) feeding primary and secondary sound systems for active noise cancellation, voice capture, and immersive audio reproduction is paving the way for active acoustics—something that the audio industry has been dreaming of for a very long time.

I recently attended an online session promoted by E-Coustic Systems, sharing details about the company's successful installations, and I was particularly impressed with the example of the DiMenna Center Cary Hall room in NY, which was fitted with active acoustics in 2016 to serve orchestra and chorus rehearsals, and is being used extensively for orchestra recording, including for countless Disney soundtracks. This is an installation based on the original Lexicon Acoustic Reverberance Enhancement System (LARES) technology, pioneered by David Griesinger, Steve Barbar, and Neil Muncy, using digital signal processing to decorrelate loudspeaker and microphone signals in real time. Of course, after hundreds of successful installations in some of the best concert halls, opera houses, auditoriums, and even large-scale open air performances, E-Coustic Systems currently offers a perfected solution with fewer microphone and loudspeaker channels, which dramatically simplifies system architecture and features advanced electro acoustic enhancement algorithms.

And because very few homes will ever be able to accommodate the requirements of a Dolby Atmos 7.1.4 speaker setup, most consumers will only listen to a reduced, spatial audio, binaural rendered version. In the near future, at least there's a very good chance that they will be able to enjoy fully immersive multichannel and object-based sound formats inside a car. A growing number of car manufacturers are realizing those immersive sound systems will need to be there anyway—so it's a win/win.

J. Martins
Editor-in-Chief

The Team

Publisher:	KC Prescott	Associate Editor:	Shannon Becker
Controller:	Chuck Fellows	Graphics:	Grace Chen, Tori Zienka
Editor-in-Chief:	João Martins	Advertising Coordinator:	Nathaniel Black

Technical Editor: Jan Didden
Regular Contributors: Bruce Brown, Bill Christie, Bob Cordell, Joseph D'Appolito, Dimitri Danyuk, Vance Dickason, Jan Didden, Scott Dorsey, Gary Galo, Chuck Hansen, Richard Honeycutt, Charlie Hughes, Mike Klasco, David Logvin, Oliver Masciarotte, Nelson Pass, Christopher Paul, Bill Reeve, Michael Steffes, Nora Wong, Stuart Yaniger, Ethan Winer

OUR NETWORK

audioxpress | circuit cellar | LIS Loudspeaker Industry Sourcebook | Voice Coil

SUPPORTING COMPANIES

All Electronics Corp.	37	Listen, Inc.	45
Clockworks Signal Processing, LLC	53	Menlo Scientific, Ltd.	47, 57
COMSOL, Inc.	9	Meridian Audio, Ltd.	21
DSP Concepts	29	NTi Audio, Inc.	65
estron a/s	15	Oerlikon Surface Solutions AG	59
GRAS Sound & Vibration	33	Parts Express International, Inc.	68
Hammond Manufacturing, Ltd.	3	Redco Audio	37
HEAD acoustics, Inc.	51	SB Acoustics	35
Hightek	25	Syntiant	17
Hottinger Bruel & Kjaer, Inc.	13	Triad Magnetics	67
JJR Acoustics	23	Warkwyn	59
KAB Electro-Acoustics	49	Western Electric	2
KLIPPEL GmbH	61		

NOT A SUPPORTING COMPANY YET?
Contact Peter Wostrel (audioxpress@smmarketing.us, Phone: 978-281-7708, Fax: 978-281-7706) to reserve your own ad space in the next issue of the magazine.

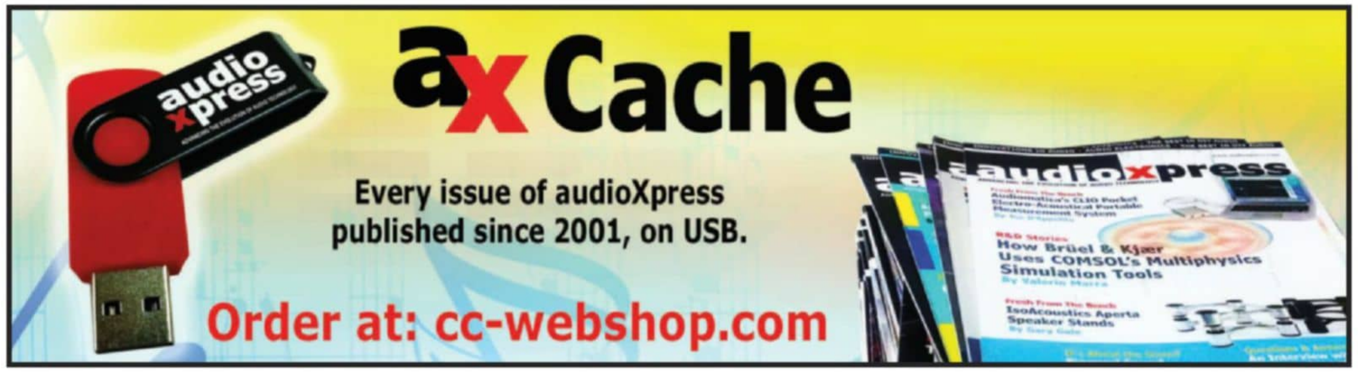

Contents

Volume 53 – No. 6 June 2022

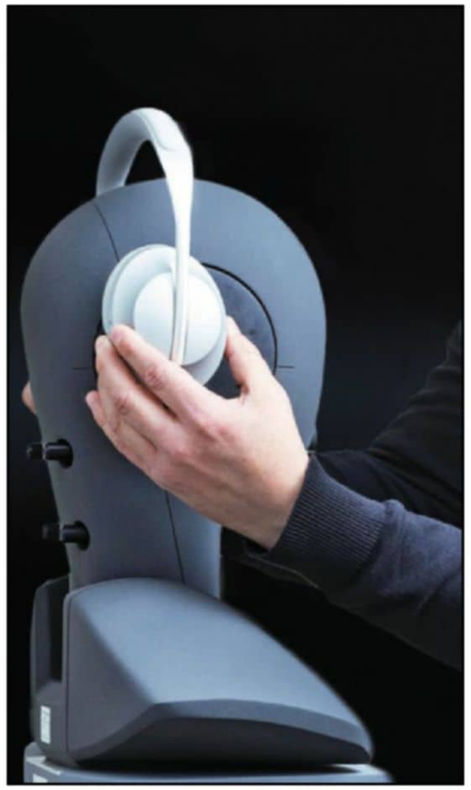

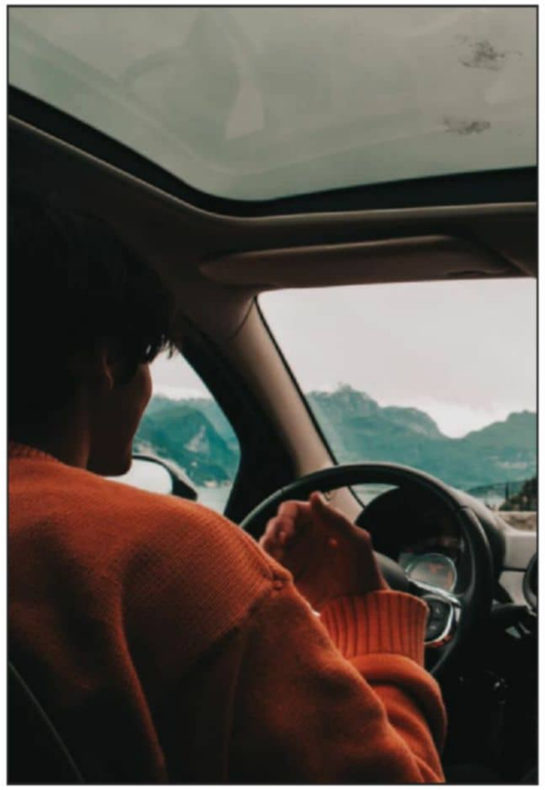

Features

42 Røde NTH-100 Headphones
Making a Late, Refined Entrance
By Brent Butterworth
With the NTH-100, microphone specialist Røde launches its first product into the crowded studio headphone market. Brent Butterworth explores how the NTH-100 stacks up to the competition—in the studio and on the test bench.

50 How Signal Processing and Filter Bank Design Impacts Sound Quality in Hearing Aids
By Dana Helmink
This article discusses the nuances of signal processing and filter bank design, exploring the most impactful decisions in filter bank types specific to hearing aids, and how ultimately each approach helps dealing with the hearing challenges that people face.

54 ANC Headset Testing
By H. W. Gierlich
This article discusses requirements for effective performance testing of communication headsets that use active noise cancellation (ANC) technologies. It addresses test setups, testing equipment and setup conditions needed to correctly assess ANC performance.

62 Building a Guitar-Controlled Synthesizer
LFO and ADSR
By Ethan Winer
Ethan Winer's two previous articles described circuits he developed for his current project that uses an electric guitar instead of a keyboard to control an analog synthesizer. This article continues with two components that are essential for all analog synthesizers: the Low Frequency Oscillator (LFO) and the Attack Decay Sustain Release (ADSR) envelope generator.

Columns

SOUND CONTROL
38 The Acoustical Effects of Room Shapes
By Richard Honeycutt
In this article, our acoustic expert focuses on the effect of the interior shape on specific acoustical parameters, including reverberation time, lateral fraction, and echo speech.

Market Update—Automotive Audio

8 Audio for All-Electric Vehicles
By J. Martins
Electric, Connected, Autonomous. That's how we describe cars of the (near) future. And that also means that cars will have to be smart, and be able to accommodate completely new requirements for drivers and passengers. And the audio industry is working on it. This article offers a broad perspective of a sector in radical transformation.

Automotive Audio Solutions

22 Upmixing and the Future of In-Car Audio
By Emil Hemström and Rüdiger Fleischer
With increasingly silent electric vehicles, car manufacturers are adding sophisticated audio systems that support many different functions, including high-quality immersive music capabilities. Those need to be consistent for driver and passengers, regardless of the source format. Learn how digital signal processors and upmixing provide the solution.

26 Evaluation of Automotive Audio Energy Efficiency
By Philipp Paul Klose
An approach to measure energy consumption and efficiency of car audio systems is essential in an industry investing in ever-more sophisticated technologies and applications. This article proposes practical approaches, including standardized test procedures.

30 An Interview with John Whitecar
A Decades-Long Journey to the Future of Automotive Audio
By Adam Levenson
John Whitecar is the VP of Product Management, Automotive at DSP Concepts. A leading expert in automotive technology, respected around the world, he believes we are only seeing the tip of the iceberg where audio signal processing technology can be applied.

Departments

4 From the Editor's Desk

5 Client Index

Websites

www.audioxpress.com
www.voicecoilmagazine.com
www.cc-webshop.com
www.loudspeakerindustrysourcebook.com
www.circuitcellar.com
www.linuxgizmos.com

@audioxp_editor
audioxpresscommunity
linkedin.com/company/audioxpress
audioxpress

COVER ART
Car Image:
Rivian R1S

Audio for All-Electrical Vehicles

BMW iX interior (Image courtesy of Bowers & Wilkins)

By
J. Martins
(Editor-in-Chief)

Electric, Connected, Autonomous. That's how we describe cars of the (near) future. And that also means that cars will have to be smart, and be able to accommodate completely new requirements for drivers and passengers. And the audio industry is working on it. This article offers a broad perspective of a sector in radical transformation.

An overview of the automotive audio market in 2022 needs to reflect the dominant trends that sweep the automotive (and mobility) industry in general. We all know that the automotive industry is going through a radical transformation, as the global pandemic accelerated technology transitions—and changes in business models—that were already under way.

The main transition that was accelerated in the past years—for many reasons, including the urgency from climate change and consequent global regulatory frameworks—are the changes in the automobile powertrain or power source. From the internal combustion engine and hybrid drive trains combining combustion engine and electric motor, to purely electric vehicles that depend completely on battery packs, all systems in the car are now more interlinked than ever. This means that until we also have hydrogen fuel cells or a secondary source of power from the one that powers the main propulsion engine, all automotive systems are basically driven by the same power source. This is motivating a completely different mindset in automotive systems design, where the interdependency of all integrated systems and power management becomes critical.

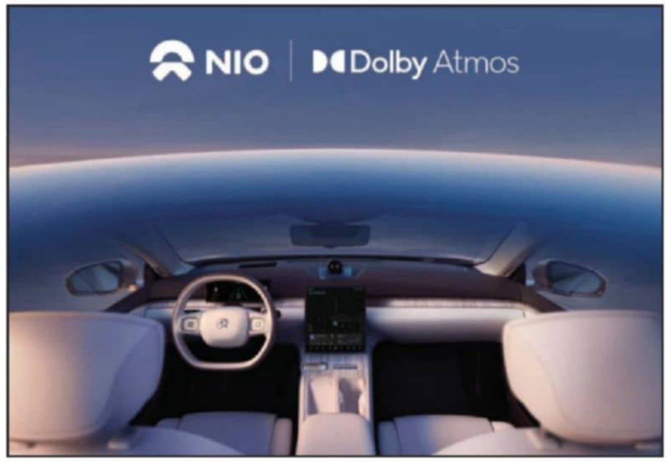

NIO, a smart electric vehicle company in China, is offering Dolby Atmos as standard on all its NIO ET7s, flagship electric sedans. The vehicle comes with a 7.1.4 immersive sound system—23 speakers and a total output of 1000W—and the ability to decode and render Dolby Atmos sources, including from music streaming services.

While in traditional combustion engines and even hybrid vehicles we had a powertrain that could always recharge a battery, from where "non-critical" in-vehicle infotainment systems could be driven, modern all-electrical designs, given their reliance on a single source, define priorities for access to power in different ways. And that includes the fact that audio systems are now being used not only for entertainment of the occupants but to generate mandatory acoustic vehicle alerting systems (AVAS), and important driver assistance features, just as an example. Also different are the systems that drive and control audio systems, predominantly managed by connected automotive technology platforms, which are increasingly software-based and interdependent.

Until now, auto makers have been focused on the autonomy and efficiency of fully electric vehicles, knowing that consumers remain concerned about the approachable range on a single charge as a key metric for such cars. And no car maker wants to admit that your fully immersive audio experience and road noise reduction—not to mention the air conditioning—could bring down that range from the current average of 200 miles (in real-world driving conditions).

Probably, one of the key technology trends in the next few years will result from the need to create self-sustained systems within the car that don't draw power directly from the main source used for propulsion. As an example, while solar cells will remain insufficient to recharge the engine batteries—in part due to the fact that the car surface is insufficient for the number of cells required—they can have a key role in powering other systems, including audio. Also kinetic and other energy harvesting will be key to create non-dependent and probably self-sustaining systems in the vehicle. It's an equation of how many batteries for those systems are required,

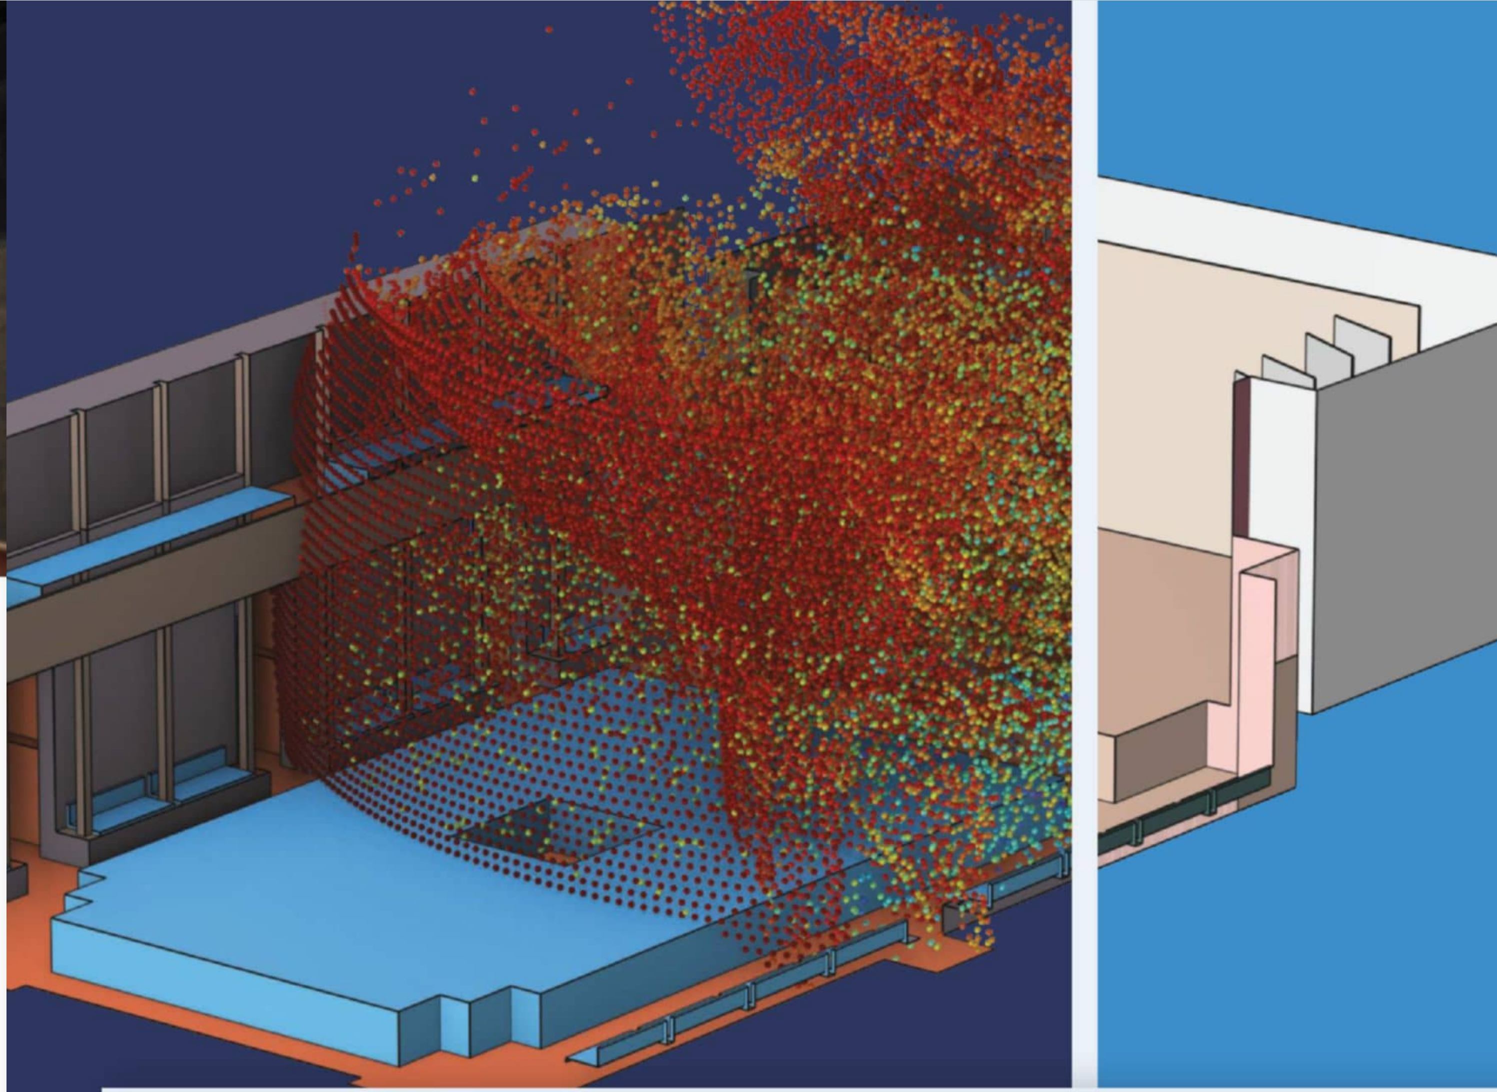

Take the Lead in Acoustics
with COMSOL Multiphysics®

Multiphysics simulation drives acoustics innovation by providing insight into every design aspect that influences product performance. The ability to account for coupled physics phenomena lets you predict, optimize, and virtually test a design under real-world conditions – even before a first prototype is built.

» comsol.com/feature/acoustics-innovation

COMSOL

The Stellantis group is the latest example of a large conglomerate of brands that embraced the fully integrated, cloud-connected platforms strategy, announcing adoption of Qualcomm's Snapdragon Digital Chassis to support enhanced safety and immersive digital experiences in next-generation vehicles, beginning in 2024.

Harman is increasingly focused on the automotive market and offers a full suite of connected technologies to the automotive and audio industries. Harman Ready Together is just the latest name for its integrated entertainment and communication management solution for automotive brands.

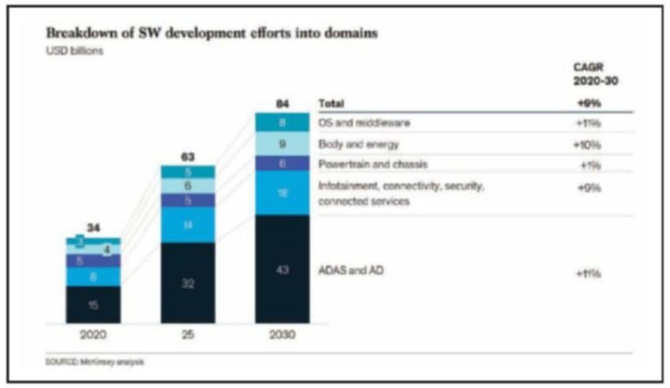

According to McKinsey (2019), the global automotive software market will expand gradually, starting from the OS and middleware levels that enable hardware/software separation, and the software functions running on ECUs, DCUs, or smart sensors (incl. infotainment), to become one of the strategic development areas for the automotive sector.

how much they weigh, and how that impacts the car performance and range.

In the medium and long term, it is a very important trend that needs to be considered by anyone designing automotive systems. The technology-specifics of this are also out of the scope of this audio-centric publication, but I will just mention the fact that current systems for fully-electric vehicles are being designed to be powered at 24V or 48V, finally leaving behind the barriers that for so many years were raised to change from 12V. This, of course, was not due to the audio systems, but instead mainly because of the battery-powered rails that were designed to sustain the connected systems within the vehicle—and audio is just a part of that. In 2019, the industry was already pushing 48V systems as it was visible from power electronic component manufacturers' announcements, but in 2022, higher voltages are part of the basic requirements in all new designs.

The shift to 48V seems to be a logical progression to sustain infotainment systems, climate-control, driver-assistance, and connectivity, as it is expected from electric vehicles—even before we consider full autonomy requirements. But the main reason for the transition was because 48V allows higher currents and minimizes resistive losses, boosting the overall efficiency of the drivetrain components (not to mention improved regulation, lower losses, and less fire-risks). The automotive audio industry manufacturers and system integrators obviously appreciate and applaud the move.

Suspended in Time

In 2020, all major market research firms familiar with the automotive industry were predicting slower growth in 2021 and 2022—but still growth. But we need to separate those projections from vehicle output, vehicle sales, and overall business volumes. Even while selling less vehicles due to the combination of factors triggered by the global COVID-19 pandemic, the industry is seeing record business based on a combination of price increases, growing investments and the sheer number of new players involved.

As the Consumer Technology Association (CTA) reported, the automotive sector significantly expanded as reflected by the number of CES 2022 exhibitors and brands that signed up for the show. Even with the uncertainties of attendance, more than 175 companies (a 12% increase from CES 2020) joined the transportation and vehicle technology section of the consumer technology show, taking hold of the new Las Vegas Convention Center (LVCC) West Hall, and featuring top brands including Audi, Daimler AG, General Motors, and more.

Even among the startups from around the world that exhibited at the CES 2022 Eureka Park, many were registered as being active in the automotive segment. And of course, so did large technology companies including AMD, Nvidia, IBM, Intel, LG Electronics, Panasonic, Qualcomm, and Samsung Electronics.

The most recent projections of growth in the car audio market beyond 2022 all mention key growth vectors from advanced driver assistance system (ADAS), voice assistants and connectivity, and growing demand for premium audio systems with noise control, all of which have implications to the overall audio industry.

While there's every reason to be excited with the innovations and developments happening fast in the automotive industry, the problem in 2022 is mainly the unpredictability of the overall supply chain and market demand. Since the beginning of the pandemic and all the major disruptions to the global economy that ensued, market research firms have been struggling to update their market intelligence and sustain their early projections and trend predictions. *audioXpress* does not conduct market surveys and vastly depends upon the data that is made available by the research firms, all claiming to be unable to predict anything consistently, due to "a perfect storm of unprecedented circumstances."

What we can confirm is that the very few projections those companies are risking to make public are basically just progressions on their last solid metrics (pre-COVID), underlined with many words of caution, near-term risks, and widened error margins explained in footnotes. If it was important to say that car sales decreased by 35% to 40% in 2020, the fact is, currently no manufacturers are able to balance their production output with the rising demand, making those numbers meaningless over a threeyear timeframe.

According to IHS Markit, persistent supply chain challenges are without precedent. "The highly synchronized global supply chain system developed over the past 30 years is under strain like never before and resolving the disruption will be less a 'sprint' and more of a 'marathon' that runs well into 2022," the company says in its latest report entitled "The Great Supply Chain Disruption: Why it Continues in 2022."

"What is unfolding in supply chains globally is not only disruptive, it is also historic," says Daniel Yergin, vice chairman, IHS Markit and editor of the report. "Moreover, the intense new focus on inflation adds to the urgency to understand what is ahead for supply chains in 2022."

And the pandemic is not the only factor. Substantial capacity, logistical and labor challenges also exist beyond the pandemic, the report says. Spotlighting the automotive sector, Matteo Fini, vice president, automotive supply chain and technology, for IHS Markit estimates that the global semiconductor and electrical steel shortages will continue, forcing automakers to limit production and pivot away from longtime assumptions such as lean inventories and just-in-time manufacturing. "The supply chain issues in automotive are unprecedented. If the question is whether this gets fixed right away, the answer is no," says Fini.

Great Expectations

So, without relying on any specific data, what we can determine are the major trends, starting with the fact that the demand for vehicles is real, that consumers expect their new-generation vehicles to be exciting and offer completely new connectivity features and expect in-vehicle infotainment systems to be fully-capable of supporting voice assistants, voice commands and offer the best possible audio quality. In summary, an environment that opens an opportunity for new companies to enter the automotive segment, while inviting a complete reboot in the approach adopted for new vehicle designs.

While for Advanced Driver Assistance Systems (ADAS), consumers still rely on the car manufacturer itself to fully warrant the features, for audio and infotainment systems, established brands are viewed as reassuring and valuable.

The global automotive industry is seeing an increasing number

Fraunhofer's automotive audio technology Sonamic Panorama is integrated into the BeatsAudio sound system of the new Seat Leon model and also provides surround sound in other new models from the VW Group. The Sonamic Panorama surround sound algorithm extracts the individual sound sources from a stereo recording, and splits them into segments—taking into account the existing loudspeakers and their positions—to generate a soundfield that is evenly distributed around the vehicle.

Many audio companies have created a close-relationship with car makers, becoming a brand-extension of each new model. For the past 10 years, Meridian Audio designed the audio systems for Range Rover vehicles, supplying the company's proprietary digital signal processing (DSP) technologies.

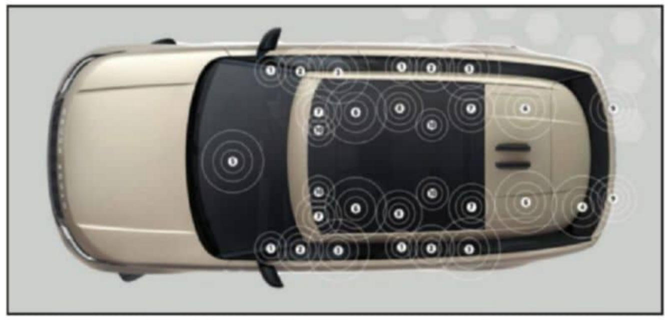

The new 2022 Range Rover features three new levels of audio performance from Meridian plus a unique Tailgate Event Suite for musical enjoyment outside the vehicle. The three Meridian audio systems offered in the new Range Rovers are the Meridian Signature System, Meridian 3D Surround Sound and Meridian Sound System. The 35 loudspeaker Meridian Signature Sound System includes new headrest speakers for the first time. These generate a totally immersive surround sound experience, creating an audio cocoon that completely envelops the passenger. New height speakers in rows 1 and 2 of the vehicle complement the surround speakers, creating natural immersion with three-dimensional spaciousness alongside new headrest and seatback speakers for rear passengers.

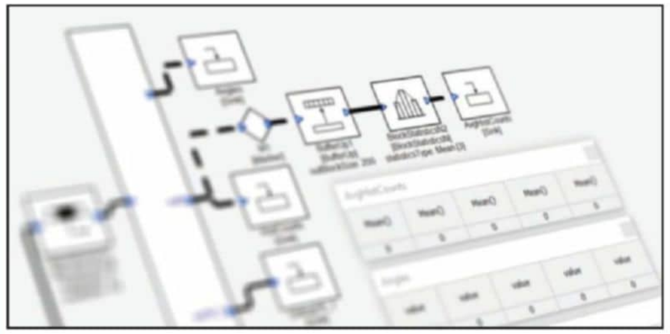

DSP Concepts Audio Weaver offers more than 550 signal processing blocks, referred to as modules, including custom automotive IP with features such as in-cabin voice communication, audio upmixing, engine harmonic cancellation, and road noise control.

of OEM announcements regarding electrification ambitions for the coming 5 to 15 years, as electric vehicles are quickly becoming the core strategy, pressured by governments and regulators. And at the same time, that industry is now being challenged with a growing number of incumbents in the electric space, including completely new companies from all corners of the globe, while the previously established Tier One supplier networks are also challenged and opening up to new system integrators and greater diversity of component suppliers.

This is also causing a major shift in the volumes that can be predicted, even by the larger automakers. While Tesla is reported to produce more than 3000 vehicles a day in all its factories, a smaller

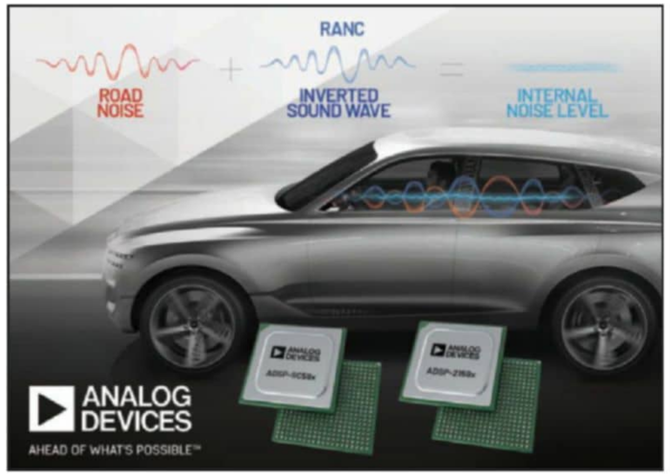

The ADSP-2156x series of processors from Analog Devices are designed to provide immersive audio in automotive and consumer/pro-audio applications, featuring a SHARC+ DSP core, upgraded from its predecessor SHARC core to natively support double-precision floating-point operations. The ADSP-2159x series adds a second SHARC+ DSP core to double the performance in a pin-compatible package, thus providing performance scalability from 400 MIPS to 2000 MIPS in a single PCB design. Finally, the ADSP-SC59x series adds an Arm Cortex-A5 core to support system connectivity in addition to the DSP performance, with models featuring either one or two SHARC+ DSP cores.

electric manufacturer such as Rivian is predicting 25,000 vehicles sent to the market by the end of 2022 (less than 100 vehicles per working day). In contrast, Toyota is now putting out less than 9 million vehicles a year, way down from its 10.7 million record in 2019. And that has nothing to do with demand, which rose sharply as soon as the pandemic effects started to fade away.

The automotive industry is facing a combination of rising costs—including from greater investments in research and development activities and hiring new suppliers in advanced technologies—greatly affected by supply chain shortages, greater logistical challenges, transportation and insurance costs, which have combined in 2022 into escalating inflation.

This is particularly significant since, according to McKinsey, the automotive industry is increasing its dependency on electronic components suppliers, accounting for around one-third of production costs and predicted to reach 50% by 2030. Those McKinsey predictions, published in 2019, have probably accelerated, given that electronic components and chips are precisely at the center of supply chain disruptions and cost increases. As McKinsey predicted, software is also becoming increasingly critical in the automotive industry. "The days of OEMs comprehensively defining specifications and suppliers delivering on them may be nearing an end. Neither OEMs nor traditional suppliers are in a position to fully define the technology requirements of new systems. Co-development between OEMs and suppliers is expected to become not just prevalent but necessary. In addition, tech-native companies are expected to more boldly enter the space."

These transversal trends in the industry, are obviously even more relevant when analyzed from an automotive audio perspective. In-car infotainment systems increasingly depend upon a number of hardware and software platforms, which the large technology players are all trying to combine. And the 2022 announcements in that area are not reassuring for audio companies that depend upon an open development environment, as heavyweights such as Harman, Qualcomm, Nvidia, and Google put pressure on the larger car manufacturers to embrace a complete ecosystem adoption approach.

Car manufacturers understand the need to keep those attempts under control and one of the obvious strategies is to keep hardware and software separate. But the new automotive design requirements increasingly invite a low-level integration of operating systems and middleware that predetermines many of the possible system options—and consequently of technology suppliers, and ultimate consumer options. As an example, Google is clearly exploring that avenue to encourage automakers to use its Android "Trojan horse" OS beyond infotainment systems, leveraging the fact that many automotive computing systems run Linux.

On the hardware front, companies such as Qualcomm are determined to offer a comprehensive platform. The Snapdragon Digital Chassis is a Qualcomm set of cloud-connected platforms with a unified architecture to support everything from audio and safety systems, to higher-level driver assistance tools, that can be updated throughout the lifetime of next-generation vehicles. While automakers have the option to adopt any or all of the Snapdragon Digital Chassis platforms, the choices of any software or systems

For automotive audio system design, the joint optimization effort by DSP Concepts and Analog Devices to support Audio Weaver AWE Core libraries running on SHARC+ processor cores streamlines the development process of feature-rich designs.

that could be adopted is clearly predetermined and filtered by Qualcomm.

Similar to what Harman has been proposing, Qualcomm also offers an integrated application framework for connected vehicle platforms, with its Snapdragon Car-to-Cloud Services—Connectivity-as-a-Service—that allows it to control with whom they decide to collaborate to support out-of-the-box connectivity. As with any system that runs Google Android, the Qualcomm platform predetermines the choices for developers and the level of features, content, and services that will be potentially "allowed"—independently of the automakers' intentions. Essentially a "gated" environment for audio technology companies. And such integrated platforms are difficult to circumvent, given that Qualcomm is a major supplier of key technologies for telematics and connectivity, including cellular radios and Wi-Fi chipsets.

Basically, for consumers, this means that their choices of a car brand (never mind the model) in the future would become more like a long-term expensive subscription of a set of services that ultimately even determines which personal devices consumers can buy and use inside the vehicle, from Dolby Atmos support to their choice of email and music streaming service.

Swedish digital audio processing specialist Dirac, a leading company in automobile audio optimization announced the new Dirac Intelligent Audio Platform (IAP), a software suite for improving automotive sound without costly hardware upgrades.

The recent announcement from Stellantis—which previously embraced Google Android for its infotainment solutions—of a multi-year technology collaboration to adopt the latest Qualcomm Snapdragon Digital Chassis platform to power the in-vehicle experiences of millions of vehicles across its 14 automotive brands beginning in 2024, is a clear indicator that the choice of car in the near future will have very broad implications for consumers—some of which they will not be aware before they sign up.

In the announcement, Stellantis says that the adoption of Snapdragon Cockpit Platforms and 5G capabilities from Qualcomm facilitates its "plan to merge all software domains" across "all major vehicle domains." The automaker group is also not shy of stating their intention to take "the driver-vehicle relationship to the next level," meaning engaging with pre-selected partners such as Amazon for services and content.

Similarly, Harman, now a wholly-owned subsidiary of Samsung Electronics is exploring the integrated approach, connecting hardware, software, and services for connected vehicles. During its recent Harman Explore 2022 event, Harman unveiled solutions that use advanced acoustic technology to streamline and enhance in-vehicle communication and entertainment. Those include "Software Enabled Branded Audio" as an on-demand software platform that upgrades the sound quality and feature set of existing unbranded car audio systems with the touch of a button.

For Harman, the approach is to offer a fully integrated systems approach that enables the consumer to activate features, from immersive audio reproduction (read optional Dolby Atmos), to acoustic environments, advanced noise management, voice assistants and communication, or customized user profiles.

The vision is also to enable shared or personalized entertainment in the vehicle, exploring advanced sound zone technology so that each occupant can enjoy their own media—an approach that continues to be somewhat elusive, even as related R&D efforts continue to increase. For those personalized experiences, Harman continues to perfect its Personal Audio Headrest solutions, with a modular design that enables OEMs to establish their preferred seating system and in-cabin integration.

A similar approach that was long explored by Bose, which was a pioneer in headrest audio systems, and continues to have the advantage that a Bose car system can also be branded Bose, while Harman needs to decide which brand—JBL, Harman-Kardon, Infinity, Mark Levinson, AKG, Lexicon—will be most valued by a consumer depending on the region of the globe or social status.

Software on Any Hardware

Aware of these potentially exciting developments for audio applications in the car, but increasingly challenging environments due to this expansive platform approach by the large technology players, key companies in the audio and digital signal processing space are having to invest in a multitiered approach that allows them to continue to support developers and integrators, no matter the choices pre-established by automakers.

As a key supplier in audio processing for automotive audio entertainment systems and voice control, DSP Concepts has been continuously investing in make sure its Audio Weaver and other software tools (e.g., TalkTo for voice control) support all hardware DSPs, processors and hardware/OS architectures. But that challenge increased exponentially with the need to also support audio designers in crafting signature engine noises or control ambient noise monitoring in all-electric vehicles.

DSP Concept's Audio Weaver is an end-to-end audio development platform that doesn't require heavy customization, includes built-in innovative audio technologies that are continuously expanded and updated, and has a proven track record of delivering advanced systems for leading brands faster than ever before. The new challenge, nevertheless, requires DSP Concepts to make sure that its resulting code supports all the latest multi-threaded, multi-core systems with optimized CPU performance.

Audio Weaver is hardware-independent, as its AWE Core runtime libraries can be embedded on microcontroller units (MCUs), system-on-chips (SoCs), and dedicated digital signal processors (DSPs) from a wide range of silicon manufacturers. Since Audio Weaver is hardware agnostic, this allows developers to continuously reuse and refine their efforts for every vehicle development, as new hardware is introduced to the market. A recent example includes the need to expand surround sound processing to support Dolby Atmos Music content, running on the latest generation of SoCs or DSPs. Because

AKM and Silentium agreed to work together on providing a ready-to-use Active Noise Cancellation (ANC) system specifically adapted for engine noise, one of the most thought-after solutions in the automotive industry, as the transition from combustion to hybrid and fully-electric accelerates.

the designs run in Audio Weaver, they can also support basic, mid-level, and high-level systems from four to more than 30 speakers, while third-party algorithms are readily available, including for upmixing and rendering to any speaker playback configuration.

For active noise cancellation and voice-control applications, the challenges increase even more, given the diversity of possible electroacoustic components, choice of audio front-ends with multiple microphone and adaptive beamforming technology. And now, the same Audio Weaver software is being used to generate, manage and tune automotive sounds for silent engines—external (AVAS) and internal, while cabin acoustics are also being reconfigured for multi-passenger configurations.

While ARM-based platforms are becoming common to power many audio processing applications, dedicated DSP solutions are still essential and in the automotive market, a clear market leader is Analog Devices and its SHARC and SHARC+ family of DSP audio processors and SoCs. The two companies have recently collaborated to optimize the performance of Audio Weaver Core runtime libraries to take full advantage of the SHARC+ core and to leverage the powerful SHARC+ DSP FIR/IIR hardware accelerators. This optimization effort results in tangible benefits for developers of automotive audio systems and other products, when utilizing audio processors in the ADSP-2156x, ADSP-2159x, and ADSP-SC59x product families.

Another company navigating the complex automotive development landscape, Swedish digital audio processing pioneer Dirac is seeing increased use of its core tools and algorithms in new-generation EVs. To leverage its previous experiences working with automakers, Tier One suppliers and system integrators, Dirac recently created its Intelligent Audio Platform or IAP, a complete solution for automotive sound in all types of car models. The modular platform, already trusted and endorsed by Volvo, Rolls Royce, NIO, BYD, and Harman, offers patented optimization and spatialization solutions, easy tuning and measurement, and an all-new way of upmixing standard stereo content to create immersive sound experiences with Dirac Virtuo. And Dirac IAP offers a complete software suite with a unified interface for sound design and optimization.

Israel-based Silentium is an industry pioneer with an already-proven ANC software solution with highly responsive algorithms used in automotive applications. At CES 2022, the company introduced its Active Acoustics software solutions for vehicle manufacturers and automotive suppliers as an integrated approach to complete in-car sound environments.

Dirac IAP is a modular platform and is unique to the car market because it offers Dirac Opteo and Dirac Virtuo state-of-the-art sound optimization and spatialization technology, combined with measurement and tuning tools in a single, unified interface and methodology for maximizing any type of sound system layout, from entry-level to premium.

The Dirac IAP solution is available in three pre-configured tiers, including advanced solutions for ensuring perfect listening at all seat positions, powered by Dirac's patented MIMO mixed phase correction, magnitude and time domain correction technology. Automotive engineers can use the available toolsets in Dirac's IAP to meet a variety of design goals, including cabin sound optimization, center image optimization, bass enhancement, soundstage widening, sound field control, surround-sound integration, and more.

The guided process provides full support to visualize target sound performance and experience every adjustment immediately with virtualization. Prototype cars are only required for measurement and final tuning. All computer-assisted tuning is handled virtually and predictably, reducing tuning time by as much as 70%. Crucial for developers is the possibility to work with the full signal chain and third-party audio technologies. All results can be shared globally throughout an organization, and Dirac IAP is supported by the industry's most common chipsets and frameworks.

Noise and Active Acoustics

In our 2020 and 2021 Market Updates, we characterized the main technology trends in automotive audio quite comprehensively, and those haven't changed. What I believe has increased in 2022 is the importance that is being given to research in voice recognition and control and supporting systems, while at the same time there is a greater realization of the concept of active acoustics in the vehicle and its increasing contribution toward many highly desirable, overlapping applications that could leverage components from audio systems.

A great example of this is provided by Silentium from Israel, which started its efforts focused on developing unique Active Noise Control (ANC) technology and sound management solutions that quickly proved effective in vehicle interiors. As Silentium expanded its innovative algorithms that adaptively follow the changes in the noise spectrum and are able to eliminate up to 90% of a sound profile or frequency range, the company realized its approach was effective by itself for low/medium-frequency noise in other potential commercial solutions including office noise and industrial workplaces.

The challenge for Silentium in automotive applications was the need to rely on complete systems to implement its algorithms. Basically, this means the need to run software on a specific computing platform/processor and leverage existing audio systems as much as possible. The active acoustics approach, which is less dependent on number of channels and speakers in the car—and the type of speaker drivers—helps to support multiple functions as a single efficient system. As an example, it offers effective noise reduction, while improving voice intelligibility for both communications and voice control. And it also creates an ideal acoustic foundation to extract maximum performance from the audio system, with less DSP resources, when playing music in immersive formats, as an example.

Increasingly, this is encouraging audio companies to work together to design and build integrated systems. Silentium recently announced a license agreement with Asahi Kasei Microdevices (AKM) to develop engine order cancellation (EOC) solutions that reduce the impact of engine noise in the car cabin. By incorporating Silentium's proven EOC software, combined with AKM semiconductor and sensor technologies, this solution enables vehicle manufacturers to quickly implement systems that improve fuel efficiency by reducing reliance on heavy noise dampening materials.

During its 15-year history of development efforts in this field, AKM expanded its noise canceling (EOC) and main driving path (Active Road Noise Cancellation or ARNC) solutions originally developed for Japanese automakers, to create a global customer service, ready to support the automotive industry in general.

The ready-to-use solution, benefiting from a large-scale integration, continuously monitors the engine noise and RPM as a reference, while cabin microphones provide a reference to the DSP platform running Silentium's algorithms, which control

Continental continues to explore developments in its speakerless audio system Ac2ated Sound, originally developed in partnership with Sennheiser, with implementations in new concept vehicles.

the built-in sound system in the car, creating an effective noise cancellation loop.

The company has been working on several global vehicle production programs for low cost, flexible ARNC software integration—offering full support from planning to production stages. Already available in new cars from Jaguar and Land Rover, the broadband ARNC technology works across an unprecedented spectrum of frequencies (20Hz up to 1kHz).

Warwick Acoustics' electrostatic transducers. Unlike conventional speakers, electroacoustic panels can be molded into diverse form factors and incorporated into almost all interior surface designs. This creates additional space within vehicle interiors for passengers, as well as freedom for automotive manufacturers to innovate visual interior designs, creating new and distinctive in-cabin experiences. The ability to incorporate these panels in the cabin ceiling and headrests is among the most appealing for new immersive audio solutions.

Silentium is also scaling up the development and roll-out of its Active Acoustics software solutions for vehicle manufacturers and automotive suppliers. As people become more aware of in-vehicle sounds, Silentium is expanding its approach to complete in-car sound environments, reducing, cancelling or enhancing any sound inside all types of vehicle.

A key component to realize effective acoustic environments in the car are of course the types of transducers used, and we in the audio industry are clearly anticipating a shift that has been long in the works. *audioXpress* has discussed many new thin-driver topologies being explored by loudspeaker manufacturers and we have mentioned more radical approaches, such as the Continental actuator solution, Ac2ated Sound, developed in collaboration with Sennheiser, effectively creating a "speakerless" audio system for the vehicle interior that simply "excites" surfaces inside the vehicle to produce sound. No doubt, that approach promises the greatest possible weight reductions, which is a key factor in electric vehicle efficiency. The challenges remaining are related to frequency response, lack of directivity control, and overall performance, restricting applications.

On the opposite end of that approach, there are credible "thin" speaker alternatives using planar magnetic transducers. Companies active in that field have been researching the potential for automotive applications and UK-based Warwick Acoustics recently announced that it received the required financial support to promote its unique electrostatic transducers "to disrupt the $8 billion in-car audio market."

Warwick Acoustics' ultra-light and low-power patented electrostatic technology is now being pitched to deliver "unrivaled, remarkable and immersive" listening experiences to the premium in-car audio sector.

But the benefits for the company's unique electrostatic transducers in automotive applications are not restricted to very high-quality reproduction. According to Warwick Acoustics, its panels are up to 75% lighter, consume up to 75% less power, and offer significant interior design and packaging freedom. Probably the greatest obstacle is Warwick Acoustics' reliance on combining its patented electrostatic transducer with a matched Electronics Drive Module. In compensation, the panels offer automotive designers and engineers significant freedom over speaker location and cabin integration.

And as automotive manufacturers (OEMs) set key sustainability objectives, Warwick believes its panels will deliver breakthrough speaker solutions for OEMs by using 100% by mass upcycled and recyclable materials. Additionally, Warwick Acoustics' electroacoustic panels do not use any unsustainable rare earth elements whatsoever—which is increasingly a major consideration for the manufacturing process.

Warwick Acoustics says it is now able to supply these electroacoustic panels in automotive-grade solutions, ready to respond to OEMs' ongoing transition to electrification and life cycle sustainability targets.

Resources

"The Great Supply Chain Disruption: Why it continues in 2022," IHS Markit, https://ihsmarkit.com/Info/0122/great-supply-chain-disruption.html

"IHS Markit: Why the Great Supply Chain Disruption Will Continue in 2022" January 2022 https://news.ihsmarkit.com/prviewer/release_only/slug/bizwire-2022-1-20-ihs-markit-why-the-great-supply-chain-disruption-will-continue-in-2022

"Automotive Electronics Market Size, Share & Trends Analysis Report By Component (Electronic Control Unit, Sensors, Current Carrying Devices), By Application, By Sales Channel, By Region, And Segment Forecasts, 2021 - 2028," Grand View Research, May 2021, www.grandviewresearch.com/industry-analysis/automotive-electronics-industry

"Automotive Software and Electronics 2030, Mapping the Sector's Future Landscape," McKinsey & Co., www.mckinsey.com

Sources

Analog Devices, Inc. | www.analog.com

Asahi Kashei Microdevices Corp. | www.akm.com

Continental | www.continental.com

Dirac | www.dirac.com

DSP Concepts | www.dspconcepts.com

Fraunhofer Institute for Integrated Circuits IIS | www.fraunhofer-sonamic.com

Harman Automotive | www.harman.com/Automotive

Meridian Audio | www.meridian-audio.com

Panasonic | www.panasonic.com

Qualcomm | www.qualcomm.com

Silentium | www.silentium.com

Warwick Acoustics | www.warwickacoustics.com

Bose Automotive Systems

Active Sound Enhancement for the Latest GMC Hummer EV Supertruck

With the launch of the 2022 GMC Hummer EV Pickup, Bose Automotive Systems introduced its first-ever application of Electric Vehicle Sound Enhancement (EVSE) technology in the world's first all-electric supertruck.

Bose EVSE technology, along with 14 high-performance speakers, are standard on the all-new GMC Hummer EV, helping to deliver a dynamic audio response for drivers and passengers. Engineered to replace the mostly silent interior that comes with operating an electric vehicle, EVSE technology sends custom-designed propulsion sounds into the cabin through Bose speakers to refine and enhance the acoustic environment.

During the vehicle development process, Bose engineers worked closely with GMC to integrate the sound enhancement technology. Using a proprietary Bose algorithm and data gathered from the vehicle, EVSE helps mask undesirable noise produced by the electric propulsion system, while at the same generating distinct sounds for different drive modes.

In Normal driving, the sound is subtle, providing a layer of audio texture to the vehicle's on-road performance. Shifting into Terrain Mode introduces a heavier bass tone to help keep the driver connected to the vehicle's performance. As the GMC Hummer EV's torque increases, the sound gradually goes up, keeping the driver in tune with the vehicle's power and movement over various types of unpaved terrain. In Off-Road Mode, a deeper rumble is heard as the vehicle moves quickly across rough surfaces, such as dirt, gravel, and sand.

The impact of EVSE technology is most pronounced in the GMC Hummer EV's driver-selectable launch-control mode, called "Watts to Freedom," which enables the driver and passengers to feel the full power of the supertruck's acceleration from 0-60 MPH in a GM-estimated 3 seconds.

As the vehicle gears up for launch from a dead stop, EVSE technology is engaged to help prepare the occupants with audio cues. A purposefully crafted mix of electric guitar riffs and feedback gradually builds anticipation as the Bose subwoofer creates vibrations and the vehicle uses its adaptive air suspension to lower its height and center of gravity. The sound then shifts, signaling when the vehicle is ready for takeoff before the driver can finally unleash the Hummer EV's 1,000 horsepower and 11,500 lb.-ft. of torque down the road.

The sound system has been designed and engineered specifically for the unique cabin acoustics of the all-electric GMC Hummer.

For music and audio entertainment inside the GMC Hummer EV, Bose applies its 40 years of experience in the automotive industry to engineer a 14-speaker sound system that delivers deeply impactful listening. The Bose surround sound system features intense low-frequency performance, powered by 6" × 9" wide-range woofers—one in each front door—and an all-new bass box design, positioned under the rear seat and visible to passengers with a Bose logo.

The system also includes proprietary Bose Centerpoint digital signal processing for enveloping surround sound in every seat and Bose AudioPilot signal processing to preserve the listening experience in all driving conditions — even when the Hummer EV's removable roof panels are off.

www.boseautomotive.com

The new GMC Hummer EV pickup features a 14-speaker Bose surround sound system, Active Sound Management, AudioPilot, Centerpoint, and SurroundStage technologies.

Measuring Automotive Max SPL

By Steve Temme Listen, Inc.

I am currently a participant in an Audio Engineering Society (AES) technical committee working group on automotive audio. This diverse group of about a dozen worldwide experts has focused on trying to standardize the way essential attributes of complex automotive audio systems are measured across the industry. Three specific measurements have been our initial focus: Frequency Response, Max SPL, and Impulsive Distortion. The committee's proposals for measurements were presented for feedback at the AES Fall Online 2021 conference in a session titled "In-Car Acoustic Measurements." I presented our work on Max SPL Measurements, Hans Lahti (Harman) presented Frequency Response, and Stefan Irrgan (Klippel) presented Impulsive Distortion; the session was chaired by Jayant Datta. Here, I will describe our proposed method for Max SPL measurements.

Let's start with why this is important. People need to be able to compare how loud an infotainment system can play in a car—manufacturers like to quote this in specifications, and consumers enjoy bragging rights about the sound level of their car stereo.

Max SPL is defined as the maximum sound pressure level (SPL) that a car's infotainment system can reproduce inside the cabin with the windows, sunroof, and convertible top closed. There are many ways this can be measured, but to keep it simple, two different measurements are recommended—overall Max SPL and Max SPL Spectrum regardless of distortion level. The reason we don't take into account distortion when we measure the Max SPL is because it is difficult to characterize distortion in a modern-day infotainment system—these devices frequently contain much signal processing, and this makes them unsuitable for playing back the sine wave stimuli that are typically used for harmonic distortion measurements.

First, let's examine the physical test setup. Our proposed test configuration replicates the position of an average person's head in the driver's seat using a precisely and specifically positioned six-microphone array in the driver's seat. The height and the angle of the seat, the positioning of the microphones with respect to the seat, and the height and the angle of the microphones are clearly defined to ensure standardized measurements across all vehicles.

The sound system settings on the head unit—the tone control and fader—are set to the factory default setting; in most cases this is neutral or flat with no equalization. The head unit's volume control is set to its maximum level using the volume control knob or digital user interface equivalent (e.g., volume level slider).

Overall Max SPL can be measured using a microphone array with the six microphone signals power averaged by analog or digital means and connected to either a conventional or software-based sound level meter that can measure true RMS and be C-weighted, as described in the IEC-61672 standard. However, if a software-based system is used for measuring the Max SPL Spectrum, it is simpler to also measure the overall Max SPL through the software. **Figure 1** shows a test configuration that makes both measurements simultaneously using SoundCheck software, and an AmpConnect 621 audio interface.

For both the overall Max SPL and Max SPL Spectrum measurements, a broadband (20Hz to 20kHz) monophonic pink noise stimulus is used. It has a crest factor of 15dB and is played for 30 seconds to make sure the system can sustain that level continuously. This is played at

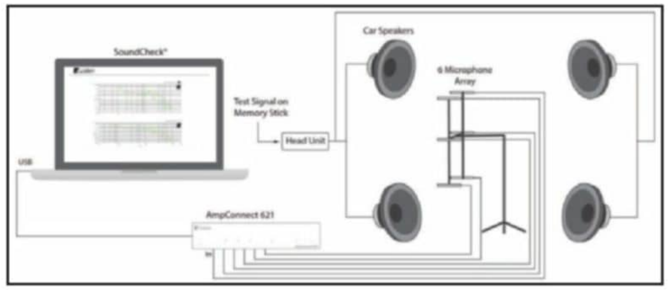

Figure 1: Test configuration for measurement of Max SPL and Max SPL Spectrum

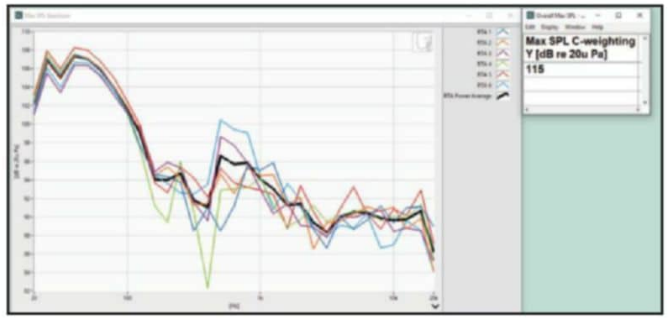

Figure 2: Max SPL Spectrum and Max SPL results

maximum volume to ensure the system is tested at the loudest signal the car will play. The sound source may come from any source—a memory stick, a CD, or Bluetooth from a smartphone or auxiliary line in. The average SPL in dB(C) is measured for 30 seconds. This is called a Leq measurement, and it takes the spatial average of the six-microphone array, power averaged, to get the overall Max SPL level (**Figure 2**).

The Max SPL Spectrum is measured using a real-time analyzer set to 1/12 octave resolution, 30 second linear averaging time and no waiting. This enables us to measure the level versus frequency irrespective of the human ear's perception. The Max SPL is recorded at each microphone simultaneously from 20Hz to 20kHz and the power average calculated (Figure 2).

Listen offers a pre-written SoundCheck test sequence that measures both the Max SPL Spectrum and a single, power averaged value for Max SPL in line with the working group's proposed guidelines. This enables consumers and manufacturers to measure the maximum overall SPL and maximum SPL versus frequency that a car's infotainment system can reproduce inside its cabin. The sequence uses the method and test configuration with a six-microphone array in either the driver or passenger seats. It takes advantage of Listen's 6-in, 2-out AmpConnect 621 audio interface, which seamlessly integrates with the software-based multichannel analyzer to measure, display, and average the results from the six microphones in real time, and power average them to calculate Max SPL. This sequence may be downloaded free of charge from Listen's website. More details about these measurements, and the other measurement proposals developed by the technical committee, will be presented at the 2022 AES International Conference on Automotive Audio, June 8-10, in Dearborn, MI.

www.listeninc.com

Upmixing and the Future of In-Car Audio

By
Emil Hemström and Rüdiger Fleischer
(Dirac)

With increasingly silent electric vehicles, car manufacturers are adding sophisticated audio systems that support many different functions, including high-quality immersive music capabilities. Those need to be consistent for driver and passengers, regardless of the source format. Digital signal processors and upmixing provide the solution.

User experience innovations popularized by electric vehicles (EVs) are raising buyer expectations with features such as semi-autonomous driving, rapid charging, quieter cabins, and enhanced in-vehicle mobile connectivity delivering greater user satisfaction. Those heightened expectations, combined with the trend of EV drivers spending time in the cabin while charging, have increased customer interest in high-quality audio technologies that deliver immersive music through the use of DSPs with advanced algorithms.

As home audio systems have added new processing capabilities and compatibility with leading surround sound and immersive audio protocols, customers have become accustomed to enjoying lifelike reproductions of virtually any recorded music at home. This new consumer standard makes it obvious when a vehicle delivers comparatively low-quality audio, causing automotive customers to feel left behind.

With a majority of music content now delivered via streaming services, publishers, producers, and content creators can easily offer multiple different channel mixes to satisfy various immersive audio standards. However, these content owners are unlikely to remix most of the hundreds of millions of existing two-channel recordings, opening the door for a technological solution to improve immersive qualities and match the surround sound listening experience provided by modern standards.

Following the upward trends in buyer expectations and technical audio capabilities, audio system designers are increasingly looking to high-quality surround sound upmixing to provide immersive audio experiences from any source through algorithmic channel spatialization that doesn't require human adjustment.

What Is Surround Sound Upmixing?

Surround sound upmixing is an advanced software feature of digital signal processors (DSPs) that digitally restructures a song's channel separation in real time, achieving truly immersive multichannel experiences regardless of the input source, musical style or original mix characteristics. The technology analyzes source content and splits

JJR ACOUSTICS, LLC
EXCELLENCE IN AUDIO

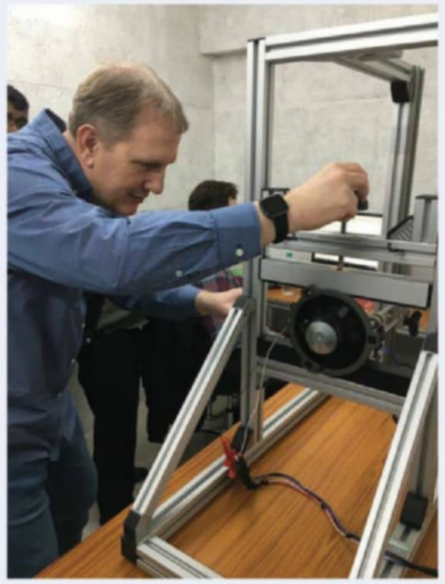

Celebrating **ELEVEN YEARS** of
Audio Consulting & Engineering Services

Specializing in advanced OEM
acoustic systems and audio products for
the world's leading manufacturers since 2011.

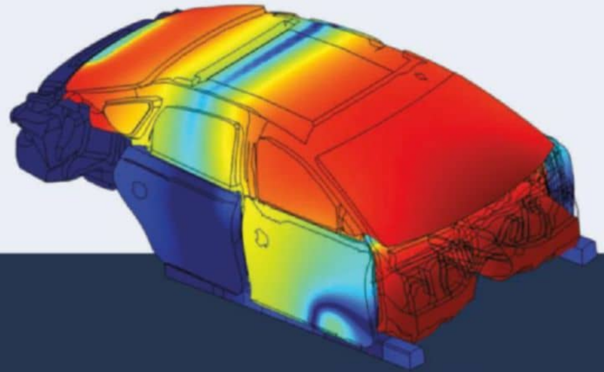

- Automotive and Consumer Audio System Design
- Automotive and Room Acoustics
- CAE and Mathematical Simulation, Optimization

- Audio and Acoustic Research & Development
- Loudspeaker Design

- OEM Business Development
- Marketing and Brand Audio Development

- Engineering Management
- Competitive Benchmarking Analysis

Contact us today to see how we can fit into your business plan!

JJR ACOUSTICS

USA | tel 317.908.6396 | Roger Shively | sales@jjracoustics.com | www.jjracoustics.com

the signal into component parts based on spatial attributes.

The upmixing technology then arranges the component parts into any number of distinct output channels, generating a multichannel version of the same content and distributing it to multichannel-capable sound systems without requiring a surround sound processor. All of this happens as a song is playing.

While some legacy upmixing solutions are known to alter a track's spectral composition and require re-tuning of connected audio systems and filters, the latest surround sound upmixing technologies employ new algorithms to minimize such effects and deliver accurate, natural, immersive playback from any input. Automotive OEMs hoping to attract discerning customers can ensure consistent premium audio experiences across a wide range of vehicles by implementing new upmixing solutions with faster, simpler setup processes that produce superior results and involve minimal stereo system re-tuning.

Surround sound upmixing can be applied to any audio signal and does not require creators or publishers to reformat or remix existing content to achieve optimal performance from various stereo designs and channel configurations. As a result, upmixing allows listeners to reinvigorate their entire music catalogs and libraries with enhanced quality and separation. Serving as a supplement to leading modern audio protocols such as Dolby Atmos, upmixing ensures an immersive listening experience for every song.

What Makes Surround Sound Upmixing an Optimal Solution?

Immersive audio has taken the audio industry by storm in recent years, with brands as big as Apple now marketing their headphones based on audio spatialization and the "real-life" listening experience it delivers. To provide equivalent quality in vehicles, OEMs require highly flexible solutions that can adapt to a range of speaker system architectures, from entry-level two-channel systems to ultra-premium systems with height speakers and subwoofers.

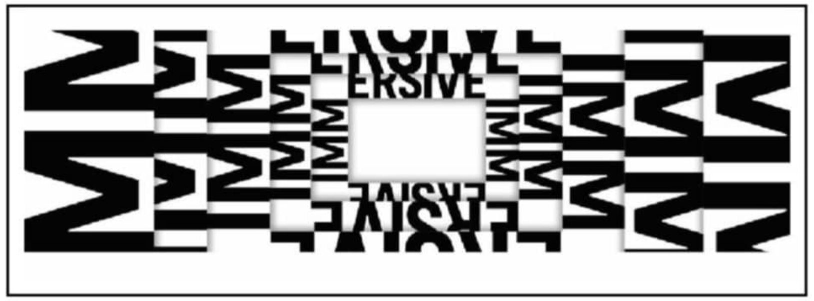

Immersive audio has taken the industry by storm in recent years.

Upmixing is the optimal solution for immersive in-car audio because it automatically works with all input sources and can be custom-tuned to maximize the effectiveness of any vehicle's speakers and stereo system. Some vehicles can be easily upgraded, such as models already using Dirac's Intelligent Audio Platform, which offers numerous patented algorithms to tune and adjust sound characteristics. For these vehicles, audio engineers can simply add our new upmixing module to the signal path without other modifications.

Not All Solutions Are Created Equal

In general, consumers are more likely to bypass or ignore a new technology if it trades one form of quality for another or suffers from performance issues. If an upmixer exhibits strange amplitude modulation, poor transient response, or any other audio degradations, audio engineers can struggle to produce satisfactory results. Low-quality upmixers can also require laborious and expensive tuning processes for each vehicle model. The latest surround sound upmixing solutions resolve both of these issues, offering much faster implementation and tuning for entire vehicle lineups thanks to the elimination of audio defects.

Competing platforms may also offer different levels of sophistication, such as the maximum number of output channels. Automotive OEMs require upmixing solutions that are highly configurable and flexible, so the same technology can be applied across all vehicles to deliver every customer a premium experience.

Automakers also want users to have control over their experiences, with various immersive options such as an "on the stage" mode or an "in the audience" mode that alter channel separation to provide distinct spatial characteristics. These types of end-user adjustments are only possible through newer, more powerful upmixers that eliminate artifacts automatically in real-time.

The options and capabilities of a solution are reliant on how much power the DSP provides, affecting the number of virtual channels, availability of presets and number of speakers that can be connected.

Vehicles of the Future

While the long-promised revolution of self-driving vehicles has yet to be fully realized, the concept itself introduces major opportunities for changes in how commuters spend their time in a vehicle. Large digital displays are already becoming commonplace, and in a self-driving vehicle, it's natural to think that users may shift toward video content, with the car's stereo functioning as a surround sound system.

A high-quality upmixing solution can turn virtually any audio source into an immersive soundscape, which is even more crucial when the listener's attention is fully focused and not split between operating the vehicle and enjoying the audio.

As vehicles continue to advance, so will the demand for premium in-car entertainment experiences, and we are hard at work leveraging Dirac's scalable technology to deliver a product roadmap with upmixing solutions for every class of vehicle.

It's All About the Experience

Automakers and their audio engineers are gaining new tools to improve the in-car audio experience, and upmixing is currently the only technology that can intelligently re-mix existing audio signals in real time. Finally, automakers can deliver a luxury audio experience without needing artists to personally remaster tracks for each different stereo.

With high-quality upmixing like that offered by Dirac's solution, automakers can optimize every song to deliver premium audio in virtually any vehicle, allowing users to elevate their existing music libraries and hear them in a whole new way.

About the Author

Rüdiger Fleischer is Product Manager, Automotive, at Dirac Research. For more than 30 years, Fleischer has worked in the audio industry in various disciplines, including transducer engineering, quality management, and technical sales. He joined Dirac in 2019 and leads product management for the company's automotive business.

Emil Hemström is Technical Sales Lead, Automotive, at Dirac Research. After graduating university in 2013, he began his audio career in the studio, mixing commercials for TV and cinema. It was there that he was inspired by the sound quality transformation made possible by signal processing and subsequently joined Dirac in 2019.

Evaluation of Automotive Audio Energy Efficiency

By
Philipp Paul Klose

An approach to measure energy consumption and efficiency of car audio systems is essential in an industry investing in ever-more sophisticated technologies and applications. This article proposes practical approaches, including standardized test procedures.

You are driving through the city with your favorite song streaming from Spotify. The voice assistant warns of traffic congestion ahead, received through a Vehicle2Vehicle message—then the telephone rings and you use the hands-free system to take the call, while the acoustical vehicle alert system (AVAS) alerts pedestrians of your oncoming vehicle.

Everyday scenes such as that make one thing clear: Automotive audio systems and their embedded services are becoming increasingly important. FM/AM radio, hands-free-telephony, and media playback found their way into all vehicle segments a long time ago—on top of these, there are concierge services, multichannel streaming services, AVAS, emergency call systems, and in-car communication services gaining ground in cars (just to name a few). The overall system complexity is constantly increasing, and with it also to some degree the system's energy consumption of the cars they are built into.

Premium-branded sound packages are what the customers want and they are willing to pay for them. In some luxury cars, the premium sound packages are beside performance upgrades as the most expensive extras. Increasing amounts of money are being spent by customers on cars with what seems like an almost per model year growing number of loudspeakers and amplifier channels. Twenty or more speakers are no longer a rarity. The current BMW iX offers optionally up to 30 speakers.

On the other hand, due to legal requirements and advancing developments of infotainment functions, OEMs are transferring more and more functions to audio systems. After all, this equipment increases the attractiveness of the vehicle in many ways and offers interesting sales margins to OEMs and suppliers alike.

What About Energy Consumption?

At the specification stage, automotive audio engineers don't often have energy consumption and efficiency on their top priority list. Besides traditional measurements of quiescent current and maximum power output for speakers, the test protocols provided by the International Electrotechnical Commission (IEC), Deutsches Institut für Normung e.V. (DIN), which is a German Institute for Standardization; similar to US ANSI, and the Consumer Technology Association (CTA) are not holding much in terms of measurement of efficiency and power consumption of consumer devices and complete audio systems.

Yet this is an important factor, through the million-fold scaling in the vehicle fleet. Even minimal amendments to the design of electric loads can have an enormous effect on overall energy consumption. Even if only a minimal amount is saved in one car, while these amounts might at first glance seem insignificant and super small, they can have a big impact on overall emissions of a car fleet. Considering that about 90 million cars were sold worldwide in 2021.

How can the electrical energy consumption of in-car systems be measured in a way that is repeatable, comparable, and conclusive? The Automotive Audio Energy Efficiency Test Cycle (AAEETC) provides a solution to this. This test cycle measures the power used by various audio sources (entertainment, telephone, warning sounds, etc.) with clearly defined boundary conditions and a modeled user.

Electric Cars Are Gaining Ground

There is a further development that makes looking at the electrical energy consumption of audio systems important: The advent of electric cars. With the introduction of a wide variety of partially or all-electric powertrains in form of 48V mild hybrids, serial and parallel plug-in hybrids (PHEVs), battery electric vehicles (BEV) and fuel cell/hydrogen vehicles (FCEV) will bring about a radical transformation of the automotive market in the foreseeable future.

Through the influence of government subsidies, high fuel prices, geopolitical interests and a general environmental awareness (which is more and more influencing consumers in their buying decisions in emerging and developed markets), the worldwide market share of electric cars is on the increase. The development of sustainable products is thus becoming more and more important and using a standardized process to measure the power consumption of audio systems is only one of the many enablers here needed. Product development engineers will be able to make better decisions, and therefore make a positive contribution in terms of sustainability. Measurement enables management.

Every Additional Watt Has a Negative Impact on Range

Why every watt counts with electric cars can be seen in a simple comparison with combustion engines. Fuels can be stored relatively easily in synthetic material containers. Converted into electrical dimensions, an eighth generation Volkswagen Golf (2021 model year) equipped with a standard 50-liter (13.2 US gallons) tank can carry 488kWh worth of energy (1 liter of petrol = 9.76kWh). A tank of the 2021 model of the best-selling vehicle in North America, the Ford F150 (optional equipment), can carry up to 136 liters (36 US gallons). This is equivalent to an impressive 1,327kWh.

Electric cars are very different in that matter: The largest battery in a car currently for sale (Mercedes-Benz EQS) has a capacity of approximately 110kWh gross. Although electric cars make an outstanding impression due to the greater efficiency of their powertrains, the total volume of usable energy is still substantially lower than that of vehicles with combustion engines.

In cold climates, the range is depleted through rising internal resistances in the battery and heating losses. When developing automotive electrical and electronic (E/E) components, vehicle manufacturers must pay close attention to the consumption of the electrical components and their efficiency. Every additional watt has a negative impact on range, charging time and a direct influence on the sustainability of the vehicle throughout its entire life cycle.

Saving the Annual CO_2 Emissions of a Small Car Fleet

But how does an audio system affect the energy consumption of a battery electric drive system? It is worth having a look at a few figures for the EU here provided by reliable sources.

The average consumption of the top 10 electric vehicles sold in Europe 2021 (charging loss not included) [1], [2] is 18.03kWh/100 km (3.44mi/kWh), while 11,964km [3] (7434 miles) are driven by the average European, traveling at a speed of 45km/h [4] (28mph). The generation of one kilowatt-hour (kWh) of energy in the EU results in the emission of an average of 317g (0.698lbs.) of CO_2 [5] (with a very strong falling tendency).

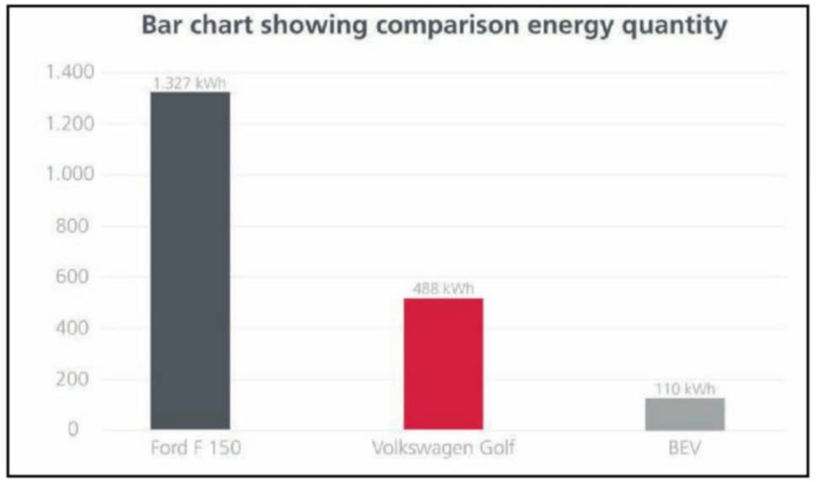

Figure 1: Comparison of carried amount of energy of popular cars and an example BEV

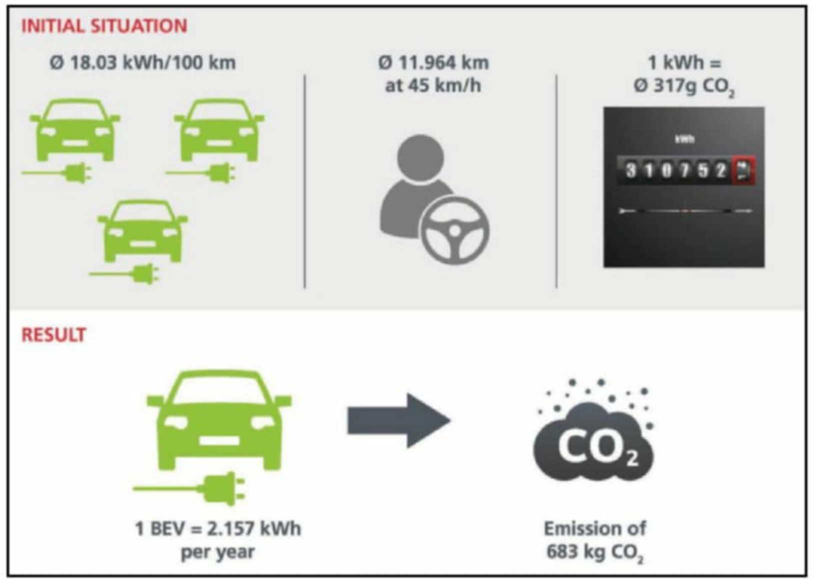

Figure 2: Environmental impact of one BEV

About the Author

Philipp Paul Klose is a senior engineer for automotive audio concepts and functions and previously served in positions such as audio system development engineer, project lead and infotainment test engineer. He holds a Bachelor of Arts (Hons.) degree in Audio Production. Within his 10 years of automotive experience, he has worked for customers such as Audi, Volkswagen, Porsche, Lamborghini and Ferrari.

Segment	Stimulus	Length
Test Preparation	- Pink noise with INT LU of -14 LUFS to calibrate the output level. Target SPL 86 dB(c) - Optional: Individual test identifier (e.g., via DTMF or voice announcement) - Optional: Measurement of max linear SPL in accordance with AES75-2022	ca. 1 minute
Start of consumption measurement		
"Music 1 (loud)"	- Pink noise with INT LU of -14 LUFS, 0 dBr attenuation	10 minutes
"Silence 1"	- No stimulus played back	2 minutes
"Music 2 (soft)"	- Pink noise with INT LU of -14 LUFS, -6 dBr attenuation	18 minutes
"Warning tones"	- Loop sequence for 1 minute: 1 kHz sinus at -20 dBFS for 2 seconds, silence for 2 seconds, 500 Hz sinus at -20 dBFS for 2 seconds, silence for 2 seconds	1 minute
"Voice Assistant"	- Loop sequence for 2 minutes: Harvard test sentences for 10 seconds, 10 seconds silence	2 minutes
"Radio 1 (music)"	- M-Noise with hard limiter, ceiling -9dBFS	10 minutes
"Silence 2"	- No stimulus	2 minutes
"Radio 2 (speech, podcast)"	- Harvard test sentences at broadcast level -23 LUFS INT	10 minutes
"Telephone"	- Loop sequence for 5 minutes: Harvard test sentences for 10 seconds, 10 seconds silence	5 min
End of consumption measurement		
Test completion	- Optional: Individual test identifier (e.g. via DTMF or voice announcement)	**Total test duration 60 minutes**

Table 1: AAEETC test description

This means that for the operation of an electric car 2,157kWh/year must be sourced (charging loss not included), leading to emissions of 683kg CO_2 a year. In 2020, a total of 1.4 million electric vehicles were sold in Europe [6]. If the electric and electronic design of these vehicles were to be altered in such a way that the continuous power usage was reduced by just 1W (-0.01235%), then this would result in a total saving of 118,090 metric tons of CO_2 per year of operation. This is equivalent to the total annual emissions produced by 173 electric vehicles or 105 average EU households! A further advantage is that, compared to the task of improving powertrain efficiency or aerodynamics, small modifications can be made to the E/E components quickly and can then be relatively easy integrated into existing series production cars.

References

[1] Sales figures by JATO Dynamics, www.jato.com/germany/tag/car-sales-germany

[2] Calculation and comparison of fuel consumption through own calculations, Spritmonitor.de, www.spritmonitor.de

[3] Odyssee-Mure, "Change in Distance Travelled by Car," www.odyssee-mure.eu/publications/efficiency-by-sector/transport/distance-travelled-by-car.html

[4] European Commission "Driving and parking patterns of European car drivers–a mobility survey," https://publications.jrc.ec.europa.eu/repository/handle/JRC77079

[5] Fraunhofer ISE, https://energy-charts.info

[6] "More than 500,000 full electric cars sold so far this year in Europe," The Guardian, www.theguardian.com/business/2020/dec/03/more-than-500000-full-electric-cars-sold-in-europe-in-10-months

Suggestion for the Automotive Audio Energy Efficiency Test Cycle (AAEETC)

How could a power-saving change like that be verified? Audio systems in vehicles are put to different use by different people: Some listen to loud music, others use the radio a lot, telephone a great deal, or perhaps use no audio sources at all when they are driving. When it came to a standardized usage test, the challenge was to take all these variables and find a mean value.

A similar challenge was also posed by the development and incorporation of the standardized test procedures for type certification for the measurement of consumption values (NEDC, WLTP, and GB T CLTC). Although the individual is not well represented by the model here, the test results allow for comparability between the competitors within a vehicle segment and meaningfulness and reproducibility between the test runs. These conditions can be transferred to the AAEETC.

Test Description of the AAEETC

In order to start the test, the tested system must be booted and ready for use. If specific filters are applied for some sources ("tunings" or EQ sets) these must be activated in the test case concerned (e.g., when the telephone is tested, the telephone filters must be active). A stabilized supply voltage of 14V to 14.4V must be provided to the DUT. To measure the current inflow, a current measurement device with an activated logging function must be used. A constant temperature of 20°C must be kept during the 60 minutes of test run (**Table 1**).

The AAEETC is a tool for developers and a suggestion of how the energy consumption of audio systems might be reproducibly and reliably measured and evaluated. Building on this, improvements on the circuit or system design can be carried out. Great potential for future improvement in system energy efficiency lies in the application of GaN-based amplifiers and multi-stage amplification circuitry. The savings expected here can easily exceed the 1W of the example calculation (**Figure 3**).

It is planned that the AAEETC test cycle will be further improved in the future through anonymously determined usage data of OEMs and providers of mobility services. The development to make the AAEETC system test a unit test for amplifier development only through the specification of a standard speaker set is also under way.

In addition, the measurement results of the AAEETC could also be used for greater transparency for the consumer in marketing materials. A fairer comparison then the "I have more watts than you" spec rivalry between competing cars could become possible.

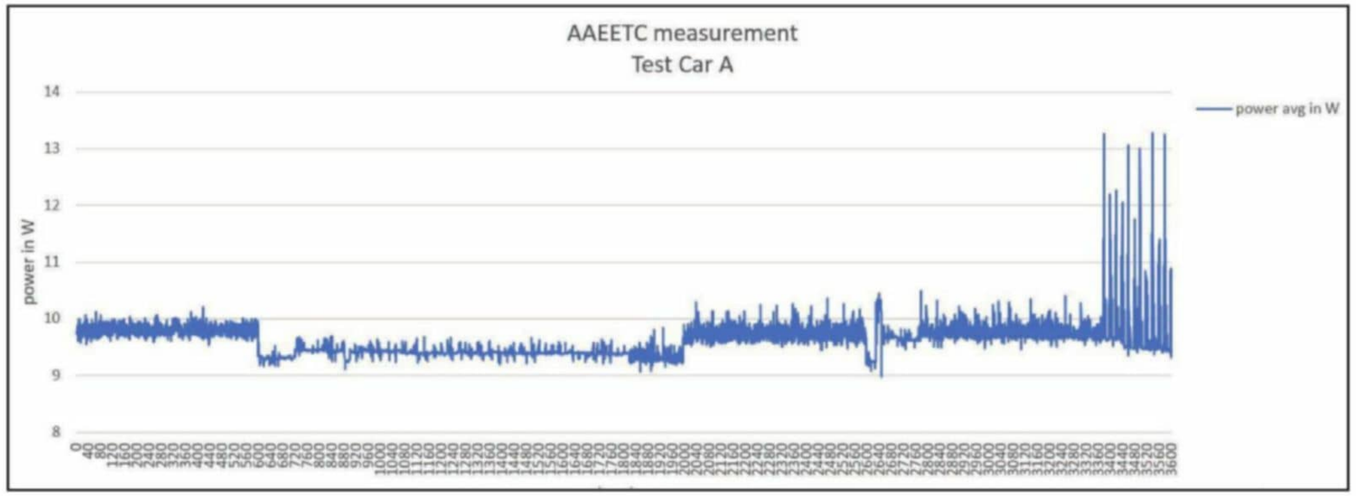

Figure 3: This graph shows the AAEETC test result of a 12-channel amplifier, 15 loudspeaker test vehicle.

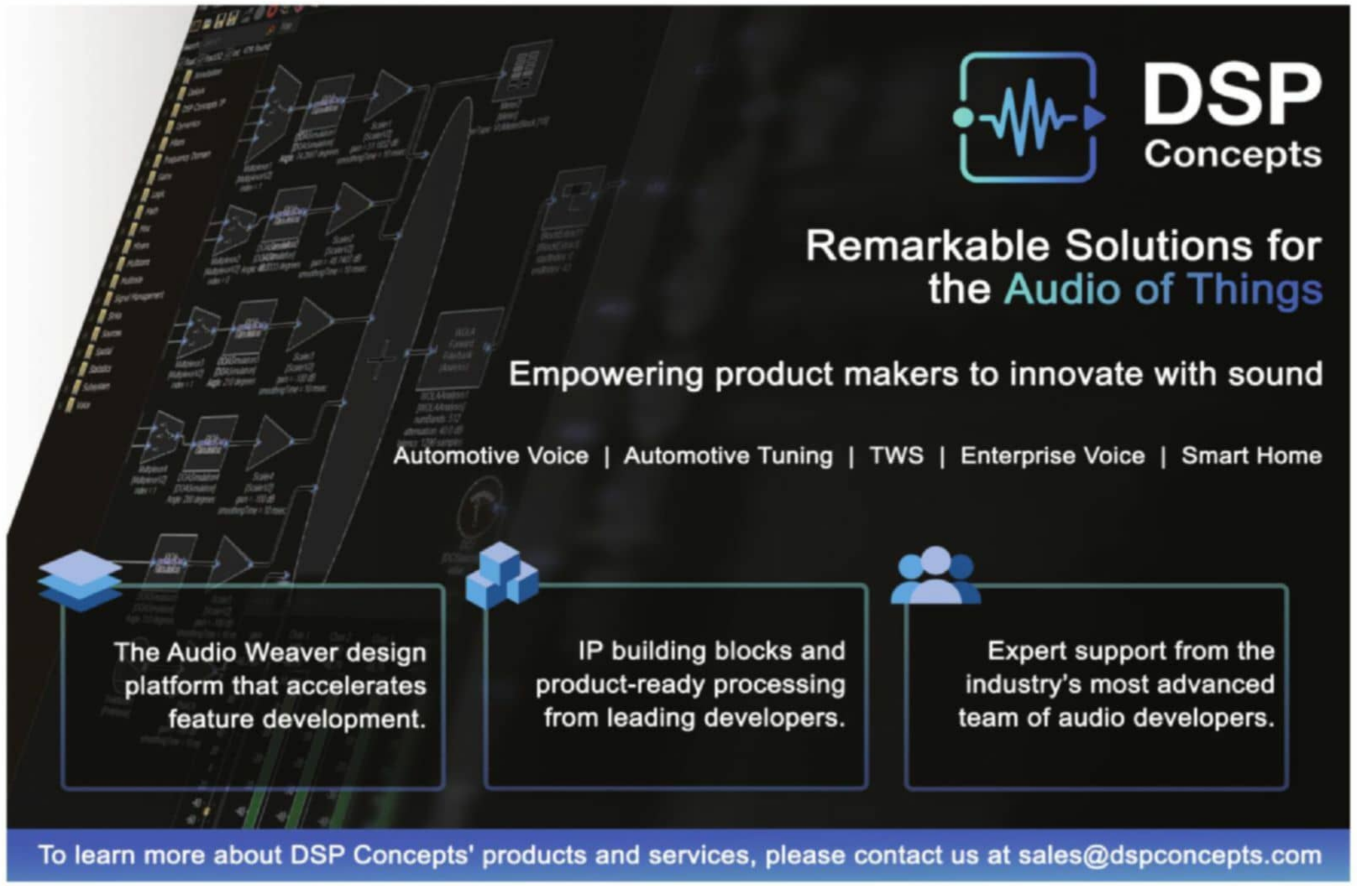

Automotive Audio Solutions

An Interview with John Whitecar

A Decades-Long Journey to the Future of Automotive Audio

By
Adam Levenson

John Whitecar is the VP of Product Management, Automotive at DSP Concepts. A leading expert in automotive technology, respected around the world, he believes we are only seeing the tip of the iceberg where audio signal processing technology can be applied.

John Whitecar's experience in automotive infotainment architecture, system design, vehicle acoustics, digital receiver, and digital signal processing (DSP) spans more than 40 years and has impacted the audio experience in millions of products. He holds 35 patents covering audio, acoustic, and receiver signal processing technologies and received the Henry Ford Technology Award from Ford Motor Co. for his work in custom mixed signal DSP-based automotive radio receiver ICs. Before joining DSP Concepts in 2017, Whitecar was Senior Manager Wireless Engineering, Principal Engineer at Tesla. We sat down with him to learn more about his inspiring career and his perspective on the evolution of automotive audio as OEMs transition toward the software defined vehicle.

Adam Levenson: You've been in pioneering roles since you started in automotive, and now you're working on groundbreaking framework technology at DSP Concepts. How did you get here? Can you share some highlights from the past, let's say, 40 years?

John Whitecar: I always wanted to be in audio. My education is in acoustics and signal processing, so that's what I studied, and I was on the fence between music and technology. My teacher said, "You'll never make money in music." But I had an interest in concert halls, reproduction systems, environments, and how it translates to the experience for the listener. How does that all work together. And when I look back through my career, everything I've been in has been about improving perceptual quality, whether it's speech quality, telephony, playback, or radio. I always asked, "How do you get the perceptual quality of the system as best as it can be?"

Levenson: I know you were in the telecom industry for the first few years of your career, in speech technology. What was your first gig in automotive?

Whitecar: I went to Ford in 1988 to work in audio system design. I designed sound systems for the Explorer—it was one of the first JBL branded systems. I did acoustic testing and helped set up their lab. So that was all about how do we create really high-end sound systems in standard mass-market vehicles. It was Explorers, T-birds, and Town Cars.

I had studied DSP previously, and then we started looking at "could we digitize all this?" To equalize cars back in the early days, you had an analog amplifier and you had to fit resistors into these circuits to get matching curves. It was weeks of effort. So I said, we can cut this down to days if we just digitize it and then take the transfer function and create the inverse transfer function. And so we invented our own tuning tools. This was the 1990s so we had 40MHz processors, and we had to get everything to crunch, and it was all assembly code. So we did the very first OEM production digital amplifier at Ford. Nobody knows about this stuff. It was all internal at Ford and they didn't advertise it—it just worked.

And then I moved into digitization of radios, and I said OK, we can take this technology into receivers. So we did a ground-up design with Texas Instruments (TI) A-to-Ds [analog to digitals] and DSPs, and came up with a whole digital signal processing baseband receiver design, and followed up with all of the audio signal playback designs. And then we found out that Phillips Semiconductor (now NXP) was off on its own doing this, and I went in and asked if we could jump on. And so we were the only non-Philips customer that was allowed to write code on their processor. So, we rewrote all the stuff that we had done, and we augmented it with the hardware capabilities that they had. Then we did a custom chip with them where we took it out even closer to the RF. So now we're digitizing right after the down converter, and then we can do all kinds of things from enhanced multipath, noise blanking, and more. The end goal was improving the perceived quality of the audio from the broadcaster.

During this era, the only source of audio in the car was the cassette deck and AM/FM broadcast. So it was all about how do we improve the perceived quality of the audio coming through the air, through the car speakers. I spent all kinds of time doing studies to determine whether one approach sounds better than another using group studies. We had a binaural setup where we could pull people in and do double-blind testing. We were pretty methodical about the whole thing. So we could make trade-offs based on the preferences coming out of the group testing. And so we did a lot of work, not only just basic down-conversion of the signals, but also on the back-end adaptation to the reception environment. And then Philips came out with Dirana, the automotive audio DSP. They moved the design to the next level with the latest Sigma Delta A-to-D technology and reduced cost. Dirana has been around for more than 20 years now.

Levenson: So it was around this time that this digitization approach became available to the wider industry.

Whitecar: Yes, and so then I moved over to Visteon and we got into the whole infotainment side. When Visteon made the decision to sell to other OEMs beyond Ford, we were looking at how we develop architectures for scalable systems. So we had to have platforms. We developed platform architectures, looked at SoC [System on Chip] development and figured out what the best approach was for each class of product, and we came up with a collection of architectures for different classes of products, and, of course, we named them after composers. I did that for quite a while, driving the whole infotainment architecture, what SoC we should go after, and how do we bring new DSP technology in.

We were doing some of the very first Bluetooth designs before there were any standards. It was the late 1990s—it was the Wild West. And then later on, the Bluetooth standard caught up, and came up with standard profiles for automotive. We had to develop the whole criteria for good voice and echo performance in a car. So we developed a series of test metrics around that, and came up with a number of corner cases for what kind of noise is encountered when we drive under certain circumstances, taking into account how that noise

John Whitecar, VP of Product Management, Automotive at DSP Concepts

The head unit in a Ford Fiesta circa 2000. "During this era, the only source of audio in the car was the cassette deck and AM/FM broadcast."

17698065

The first displays with higher-end graphics began to emerge, disrupting the market for established SoCs, such as TI's Jacinto 4 and 5.

translates over the cellular network. We spent a lot of time on that front, and then we got business with Nissan and BMW.

Levenson: All of these trends must have been leading up to the first in-vehicle infotainment systems.

Whitecar: Infotainment had its inception in the late 1980s with a display and touchscreen, but functionality was quite limited. In the 1990s OEMs were pulling together designs, but they were hardware-centric, requiring multiple ECUs. So, at this time Texas Instruments was looking for someone to define SoC architectures for the automotive market to bring system cost down by minimizing ECUs and the number of chips required to build a system. The new thing at the time was encoded audio, and the focus was creating MP3 decoders. That was the genesis of the TI Infotainment business. A whole DSP was used to perform an MP3 decode. So, I joined TI, and we pulled together requirements for a single SoC infotainment system.

Tesla was looking to develop a state-of-the-art six-tuner radio system. John said, "yeah, I can do this."

Bringing together color display, graphics, voice recognition, navigation, audio, radio, and Bluetooth connectivity into a single chip was a challenging concept at that time. It was the first time anyone had done that. We designed and pushed it out to market. The thinking, which was a bit new at TI, was to build what the market needed. Let's build this thing and they will come. They were very nervous about it, but we sold a ton of them. It turned into a $300 million annual business and we were the leaders.

By the time we got to Jacinto 4 and 5 SoCs, we had taken over much of the high-end infotainment business. And that's when Nvidia came in with very high-end graphics capabilities. We weren't really willing to go down that path because most of the fundamental chipsets came from the wireless and tablet market. Nvidia came from the gaming world with high-resolution displays and high-end graphics. TI decided not to go down that path so that's when Jacinto 6 kind of tapped out. Plenty of horsepower, but the graphics capabilities were limited. So Jacinto 6 became the second SoC with Nvidia graphics sitting on top of it.

Levenson: So you were at TI for more than a decade. What happened next?

Whitecar: Tesla called. In 2015 they needed someone to lead the broadcast radio roadmap and develop a best-in-class receiver to match the brand expectations. So I said, "Yeah, I can do this." So, we designed a state-of-the-art six-tuner radio system that was full diversity in every mode, and could seamlessly blend audio across bands. So that if you were driving in Europe and listening to a DAB station, and you drove out of DAB range, it would seamlessly blend to the FM station that was on the same channel. You could drive around the region and never have to touch the dial. That was pretty complex because it was a control system as well as a lot of signal processing to align the signals and make sure that they blended seamlessly. While at Tesla I met Chin Beckmann during orchestra rehearsals.

Levenson: Chin Beckmann, CEO and co-founder of DSP Concepts and a concert pianist!

Whitecar: Yes, she was the pianist. I was the tuba player. It was 2016. She knew I was working at Tesla, and we would talk during intermissions or during breaks at rehearsals, and she would tell me about her challenges in starting DSP Concepts. I would give her my advice, like "here's how you have to talk to the automotive OEMs," "here's how

Automotive Audio Solutions

About the Author

Adam Levenson's 30-year audio career has evolved from sound effects design, to business development, to marketing with roles as Director of Audio at Activision, GM at CRI Middleware, VP of Product Marketing at Waves, and now Head of Marketing at DSP Concepts.

you have to talk to the Tier Ones," "this is how you message to them." So finally she said, "Why don't you just come to work for us?" So I'm here, and the rest is history.

Levenson: Give us some insight into your role at DSP Concepts.

Whitecar: I'm the automotive product manager, responsible for setting the technology road maps, product road maps, understanding the market, talking with customers, and finding out where they're headed. So I'm looking at RFQs and requirements and figuring out what kind of product features customers are looking for, and then putting together a technology and product road map that aligns with where the market is headed. And so, understanding audio features, understanding the voice features, understanding things like playback features, and even other things like safety-related audio and even how audio applies to autonomous vehicles and autonomous driving situations.

Levenson: From your perspective, how has the automotive business at DSP Concepts changed since you started? Obviously, it has grown.

John Whitecar, an accomplished tuba player, and Chin Beckmann, concert pianist, met while playing with the California Pops Orchestra.

Whitecar: Yes, it's grown. At the starting point we basically created playback systems for customers, and we tuned them. It was all bespoke, highly customized. It was more like a consulting services business. Then we built out the tuning tools to help us to tune faster. In fact, one of the reasons that Chin brought me on was to help productize things, to really help scale. To go from being a great services company to packaging things up and productizing them and to create entities that salespeople can sell and that our support groups could more easily support. It was about 2017 when we made that pivot, after DSP Concepts got the A-round funding.

So Paul Beckmann (DSP Concepts CTO and co-founder) and I sat down, and we really started carving out what's needed for automotive. We started looking at how we could take some of the voice technology that DSP Concepts perfected for consumer and enterprise and bring it into automotive. It was still early for the market in automotive voice assistants. There weren't any Alexa Auto products yet, for example. And we said, "We know Alexa, we know Amazon, we know how this stuff works. Let's start promoting that," and so we started getting some bites, especially with the EVs. And that's where it started. And then the OEMs were like, while you're here, why don't you own the mics and do the telephony. So then we had playback, voice control, and telephony. That's how it started with Lucid and Rivian.

Levenson: When you think back on your vast experience, and then look at where we are now in the industry, what are the main things that have evolved?

Whitecar: So to me, there are four things that have really moved the needle. First, sensor costs have come down. Microphones used to be pretty clunky ECM modules, expensive, and you could afford to put one in the car. And then later it was like, I could put these little PDM mics all over the car. They're twenty, thirty cents, really inexpensive. Now microphone sensors are ubiquitous in cars. The other one is compute power. Compute power has exploded with these GHz [gigahertz] plus octa-core devices and quad-core DSPs. The third one is the whole EV movement and autonomous driving that has forced a totally different usage model. The car is now a moving theater. It's ultimately a different use model when we get to full autonomy and ride sharing.

Thinking back, when I was doing audio at Ford back around 1988 and then all the way to about 2005, the technologies didn't change much. In 17 years and didn't change much. What we put in Dirana or in the car DSP was sufficient. Six-channel system, 5.1, add a woofer, take away a woofer. That was really the only choice. Should I put a subwoofer or not. It was pretty simple. Dolby was trying to get in with Pro Logic, failed. DVD audio was trying to get in, failed. Super CD, failed. All of these high-end things failed because they couldn't get the customers excited about it. And so anything that was interesting in audio just didn't move because there was no inflection point that could really change the model.

Levenson: So what was that inflection point?

Whitecar: The inflection point was that voice became important. That's the fourth thing that has really moved the needle. Now with voice, the user is really interacting with the vehicle. At Visteon, we had a Jaguar with voice control in 1998. All it did was set the AC and turn on the radio because that's all you could do. It wasn't connected to the internet and it was kind of clunky, and it only worked for American native speakers. If you had any foreign accent, it wouldn't work. It was not a success.

But now, the speech technologies are fantastic. So the demand for voice technologies came out of the fact that it actually works now. Now OEMs have microphones all over the place, and that introduces all kinds of other use cases. I want to make a phone call, that's the obvious thing. I want to be able to have a conference call with everybody in the car. I want to go ask for the weather. I want to get a route. So now there's an explosion of the back-end ecosystem and use cases. That's huge. Now it's a question of which back-end do I want to use? I want to get music from Spotify. I want to get the weather from Amazon. I want to get my route from Google. I want to change the climate control, so now I just want to do that locally. So now the problem is really about how do I deal with multiple back-ends, and that's changed just in the last five years.

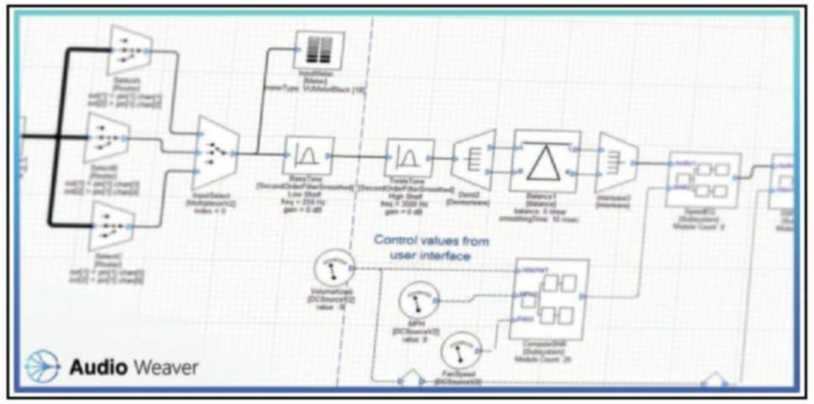

Screenshot of an automotive processing chain built in the drag-and-drop Audio Weaver Designer environment

AUTOMOTIVE

From almost no connectivity and no real good use cases to now.

Levenson: And you said that the EV movement is another thread that moved the needle.

Whitecar: So the electric vehicles brought new problems: road noise. We don't have engine noise anymore to mask the road noise. So now everybody hears everything, and it's all a trade-off of how much damping material does the OEM want to put into the car to quiet it versus some active way to do it. Given that the sensors and the speakers are installed in the car already, is there something we can just layer on with software? And that's where the whole compute thing comes in. There's tons of compute power now. We can do complicated functions in software that would have taken a completely separate ECU to do previously. And this really goes into this whole software-defined vehicle that you hear people talking about.

Back in the day, I would have a group that did nothing but design radios. I would have a group that did nothing but design audio systems and amplifiers. I had a group that dealt with the voice system. I had a group that dealt with noise vibration and harshness (NVH) and noise abatement. I had a group that dealt with safety. They were all off in their corners, designing things and making boxes for every one of these problems. And the boxes were very loosely coupled to each other through some CAN network. Now you have a system where there's like eight cores running at 1.5GHz. You can do a lot with that.

Now you're saying, "Let's centralize all of this compute power." Now all of these boxes have much more connectivity with each other, and we can do more things together, knowing about each other. The safety systems can respond to how the person in the car is using the vehicle. For example, how many people are in the car? Is the driver talking to somebody? Is the driver listening to loud music? This is all information that's available in centralized so that they can take it and apply it to all the systems. So now you have these interactions between these different audio components that were all separate before.

Now they're coming together and you've got the compute power to do it, but now you have the problem of how do you get it done? This is all software and the software modules all need to talk to each other. That's why you need a framework like Audio Weaver.

Instead of having to do an integration of an AVAS solution with an integration of a voice solution with a playback solution, these are all modules on Audio Weaver, all pulled together. Now in the Audio Weaver framework everybody is speaking the same language, all of the modules know how to communicate with each other. Now you can leverage from the synergies.

By knowing what different parts of the car are doing, we can make the system smarter. The computational piece is really what pulled everything together.

Levenson: What about internal Engine Sound Enhancement (ESE), chimes, and external Acoustic Vehicle Alerting Systems (AVAS) that come with EVs? How do you see that evolving?

Whitecar: The movement is toward creating soundscapes in the vehicle that are tied to the brand. We're talking about signature-designed orchestrated systems that work harmoniously and are branded. This is an initiative that's being worked on right now and some of the car companies actually have early versions.

Levenson: So, actually tuned with the external and internal sound of the vehicle?

Whitecar: Yes, part of the vehicle tuning. So I get into the car and I get a certain wake-up sound with chimes. The chimes all have a cohesiveness to them, they're intuitive, and there's a soundscape that's created. And it's consistent, like I know I'm in a Porsche, I know I'm in an Audi, I know I'm in a BMW. Just as iconic and identifiable as the OEM logos. This is definitely a strong trend that's going to blow up. And that's also coming out of the fact that I have more compute power and more ability to do a lot.

Back when I was at Ford, we would play .wav files, and you know we had 4KB of memory. We could go like "beep." That was it! And so it wasn't realistic. But again, the ideas have been around forever. But it's always that you have to wait for the compute power to catch up, and to have a powerful framework available to the point that now we can feasibly implement this. So a lot of times when I hear people announcing what they think are new innovations, I think, yeah, we thought of that in 1982, but we just couldn't do it. That's why sometimes I get a little cranky because it's all been thought of before!

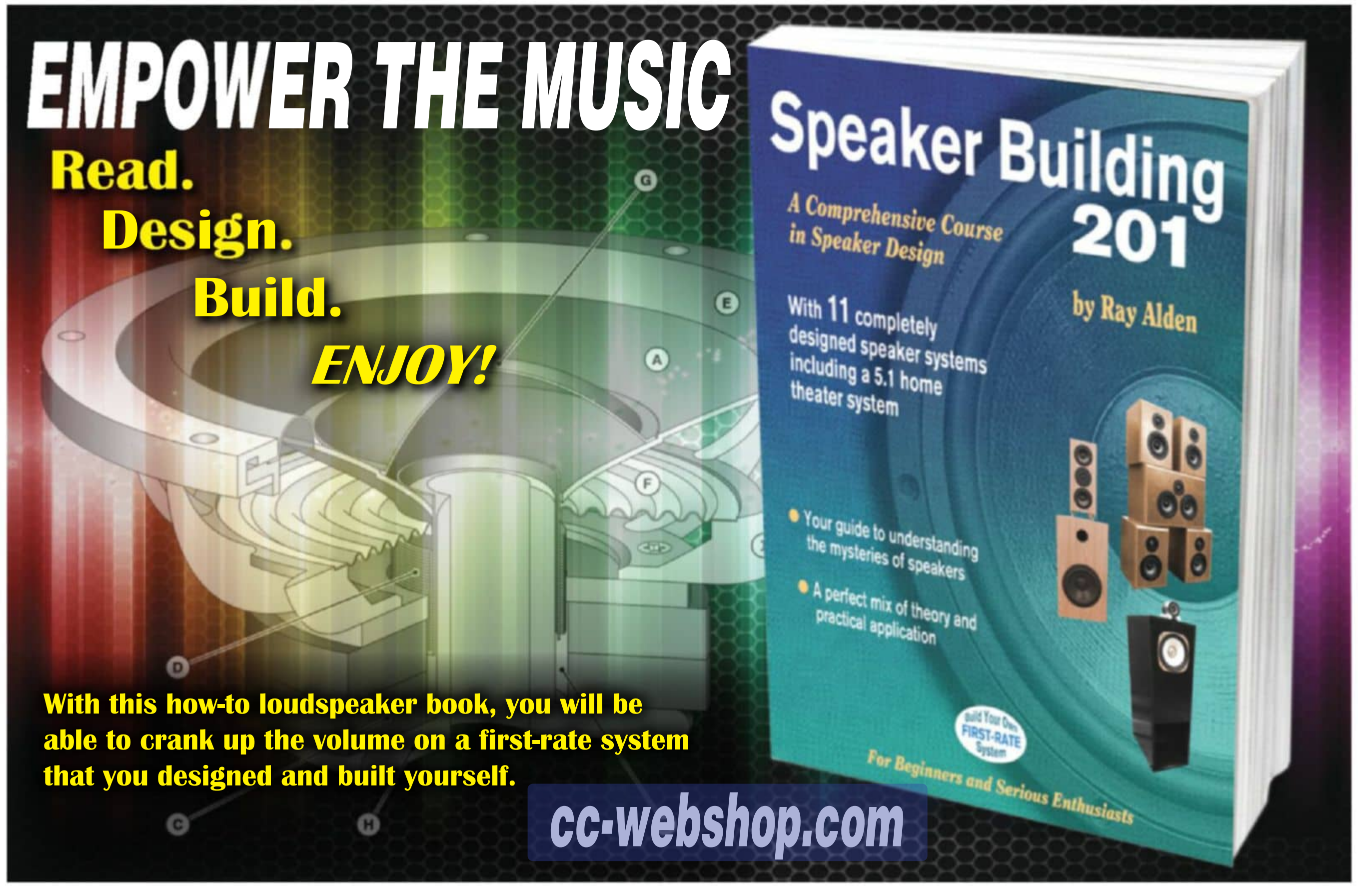

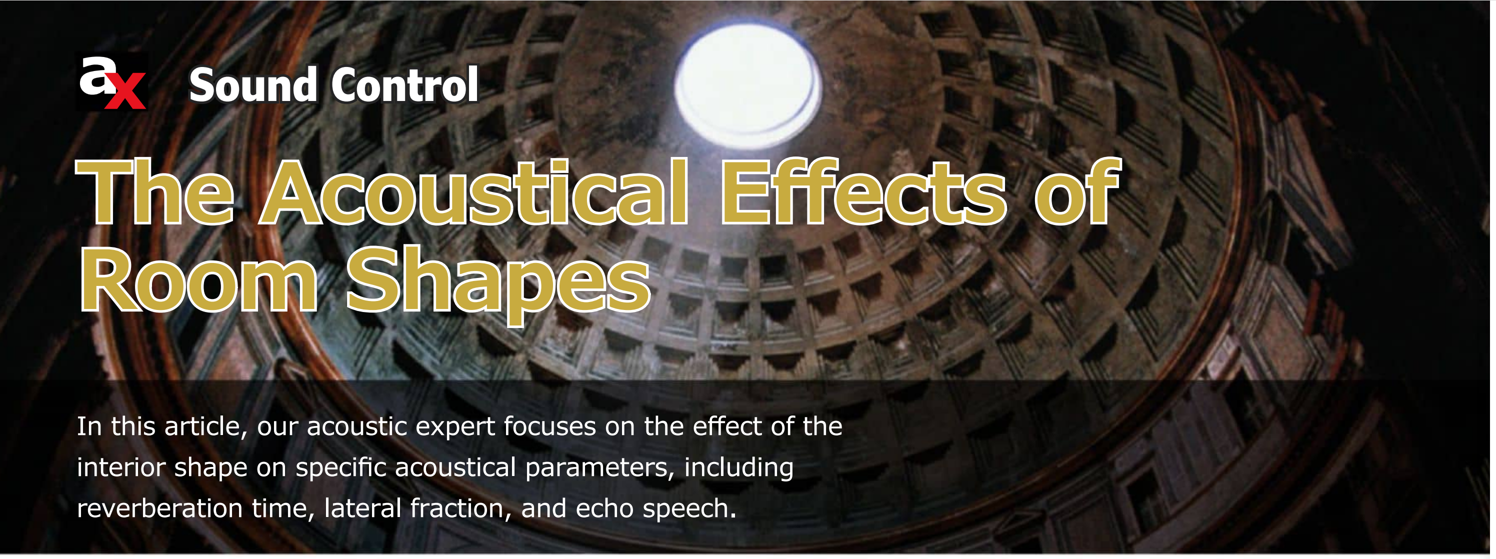

The Acoustical Effects of Room Shapes

In this article, our acoustic expert focuses on the effect of the interior shape on specific acoustical parameters, including reverberation time, lateral fraction, and echo speech.

Photo 1: The Roman Parthenon building was circular in shape.

By
Richard Honeycutt

In the world of architectural acoustics, conflicts often arise between the architect (or the interior designer) and good acoustical practice. It may have always been so. Not only were some "classic" buildings, such as the Roman Pantheon building (**Photo 1**), circular, but in the case of the Pantheon, the diameter and the interior height were the same: 141 feet. The Cathedral of Chartres (**Photo 2**) has stone walls and floor, as well as a domed ceiling whose height is approximately 121 feet (37 meters). While these ancient buildings were certainly not used as modern auditoriums are—thus having very different acoustical requirements, it is noteworthy that their use of circular elements would have created focused reflections detrimental to speech intelligibility and any music other than chants.

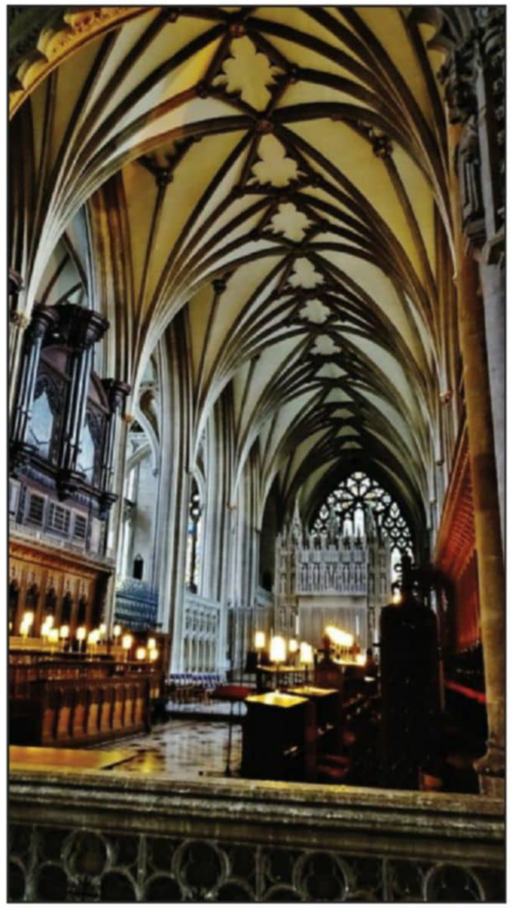

Photo 2: The Cathedral of Chartres has a high domed ceiling.

Such focused reflections not only create "loud" and "weak" regions in the parts of the room that are occupied by participants in events held within the space, but also can mar speech intelligibility because the converging reflections from, say, a domed ceiling will arrive at a listener's ears delayed by the greater distance the reflected sound has traveled compared to the direct sound.

Common Shapes and Parameters

Although Wallace Clement Sabine's spectacularly successful "shoebox" design of the Boston Symphony Hall (**Photo 3**), based upon the acoustically excellent Leipzig Gewandhaus, the Vienna Musikverein, and the Amsterdam Concertgebouw provided a clear example of how to build a fine hall for concerts and lectures, less well-conceived designs still appear today in auditoria, civic centers, and worship centers. In this article, we will concentrate on the effect of the interior shape on specific acoustical parameters.

Reverberation Time (RT) has undoubtedly been the best-known acoustical parameter for an auditorium ever since Sabine defined the concept and the means for optimizing it around the turn of the 20th century. In fact, for much of the time since Sabine's introduction of the concept of RT, it was essentially the only acoustical parameter of a hall considered susceptible to quantitative analysis and control. RT is defined as the time (in seconds) required for a sound to decay by 60dB. Research conducted later in the century revealed that audience preference is more closely related to a different but similar parameter: Early Decay Time (EDT). This is the time required for a sound to decay by 10dB from maximum. EDT controls listener perceptions of clarity and definition, and is related to the perception of intimacy, although other factors (e.g., visual) are involved as well.

Lateral Fraction (LF) is the percent of sound energy arriving from the listeners' sides compared to total sound energy. A high LF helps create a sense on envelopment in which the listener feels surrounded by the sound.

Echo Speech is a parameter defined in 1986 by L. Dietsch and W. Kraak to quantify the presence of sound reflections that listeners are likely to perceive as echoes [1]. There are two versions of the analysis: "Echo Speech" and "Echo Music" to accommodate humans' different thresholds of echo perception for the two types of program material. Any score of unity or above is considered to indicate likelihood of perceptible echo. Many consultants prefer to design spaces for echo scores under 0.8 because of possible imperfections in computer modeling.

The Effect of Different Spaces

We will consider three spaces in this article. The first is a typical rectangular school auditorium; the second, a similarly-sized auditorium having a "fan" seating plan; and the third, a near-elliptical auditorium.

Photo 4 shows a computer model of a typical "shoebox" school auditorium. The light brown structures represent audience seating: medium upholstered chairs occupied at two-thirds of maximum capacity. The purple structures are reflective acoustical ceiling "clouds" which even out the spread of sound throughout the seating area for acoustical (unamplified) sources such as actors, choirs, bands, and orchestras. The dimensions are 63'L × 71'L × 23'H. The rear wall is treated with 2" fiberglass panels from 3' to 7' above the finished floor, in order to avoid reflections from the speaker system. Additional acoustical treatment is used in order to achieve a 1kHz RT of about 1.4 seconds, which is appropriate for a multipurpose auditorium of this size. The computer model predicted the EDT and LF shown in **Photo 5** and **Photo 6**.

In order to increase the sense of visual and aural intimacy, fan-shaped auditoria are often used. **Photo 7** shows a model of a fan-shaped auditorium. **Photo 8** shows that the fan shape provides a lower EDT, and hence, we'd perceive increased clarity and intimacy. **Photo 9** shows a lower LF in the

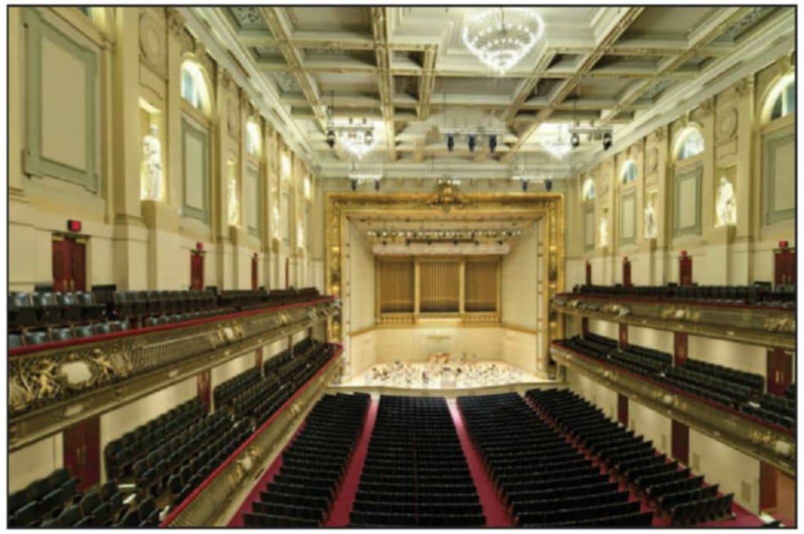

Photo 3: The rectangular Boston Symphony Hall is a fine example of acoustical excellence.

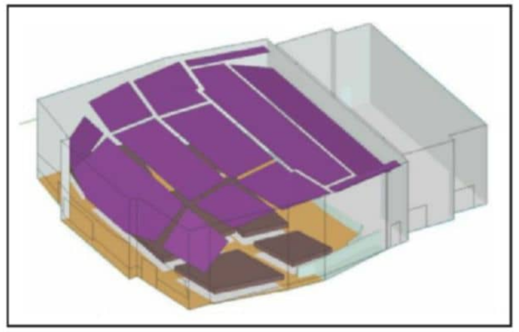

Photo 4: This typical school auditorium has a "shoebox" shape.

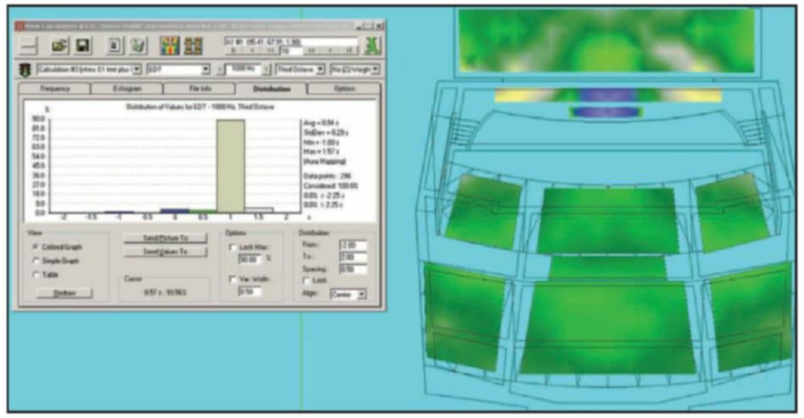

Photo 5: As this map shows, average EDT for this "shoebox" auditorium is just under 1 second.

About the Author

Dr. Richard Honeycutt fell in love with acoustics after his father brought home a copy of Leo Beranek's landmark text on the subject when Richard was in the ninth grade. Richard is a member of the North Carolina chapter of the Acoustical Society of America. Richard has his own business involving musical instruments and sound systems. He has been an active acoustics consultant since he received his PhD in electroacoustics from the Union Institute in 2004. Richard's work includes architectural acoustics, sound system design, and community noise analysis.

fan-shaped auditorium, which would result in a reduced perception of envelopment.

Sharp Focus Is for Video

Most people associate focusing with visual things: photos and videos. Focusing is usually achieved by the use of lenses, although the reflecting telescope invented by Isaac Newton in the 17th century uses curved mirrors to focus the image. The parabolic microphone is an audio analog of the reflecting telescope. Sound arriving from the front of the reflector is reflected to the diaphragm of a microphone located at the focal point of the parabola.

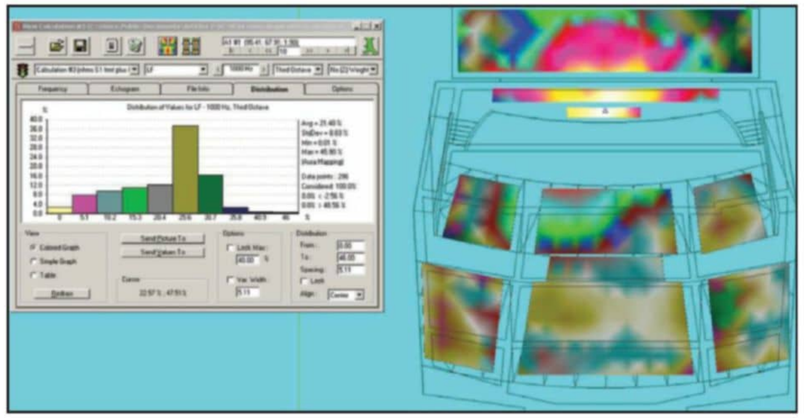

Photo 6: The LF for this shoebox auditorium averages 21.4%.

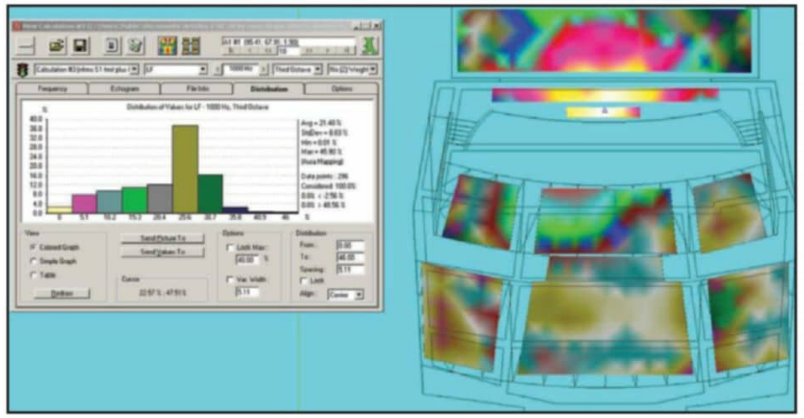

Photo 8: The EDT of the fan-shaped auditorium averages 0.79 seconds

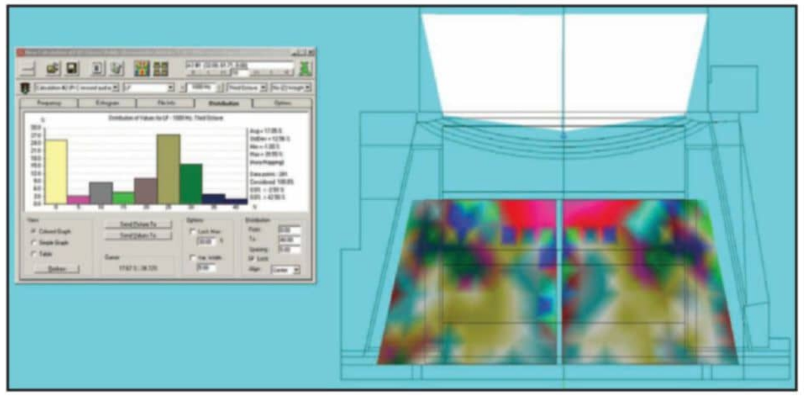

Photo 9: The LF of the fan-shaped auditorium averages 17.5%.

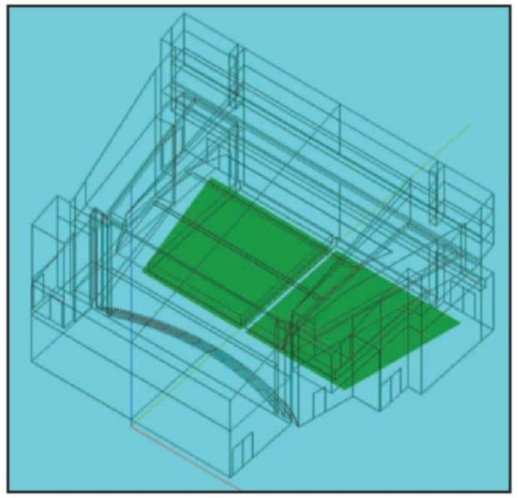

Photo 7: This computer model of a fan-shaped auditorium is viewed from the "stage" end.

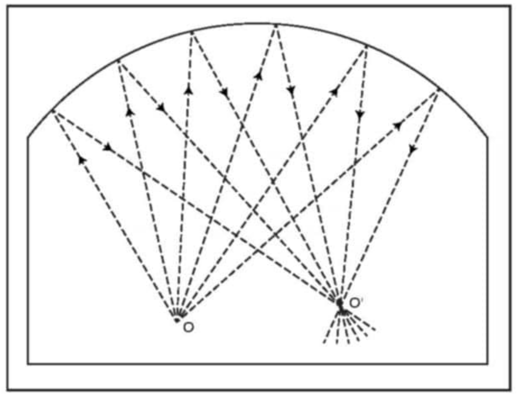

Figure 1: Conversation at point O is focused by the barrel-vaulted ceiling and can be clearly heard at point O′.

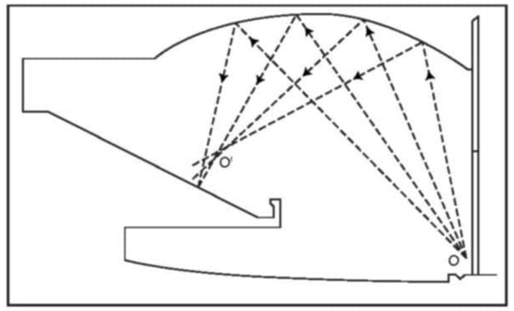

Figure 2: A domed ceiling can create detrimental echoes.

In architectural acoustics, focusing is generally undesirable, as it causes uneven sound levels at different places in a venue. It can also cause incidental noises, such as audience members having a private discussion, to interfere with program audio in lectures, dramatic performances, and concerts. In one church for which I have consulted, normal conversations in the house-left front of the nave could be clearly heard in the house-right end of the balcony due to focused reflections from the curved top of a window alcove.

Figure 1, a transverse section drawing of an auditorium having a barrel-vaulted ceiling, illustrates a similar situation. These focusing effects are reversible: conversations in the balcony or at point O′ would also be audible in the house-left front nave or at point O [2].

Although reflecting clouds such as the ones shown in Photo 4 are often incorporated in auditoria in order to increase beneficial early reflections from acoustical sources to the audience, these are carefully designed so as to not focus the reflected sound.

Figure 2 shows a longitudinal section of a reflecting ceiling that is poorly designed from an acoustical standpoint [3]. Sound from a talker just downstage of the curtain will be focused at a point about one-third back in the balcony. If the ceiling is domed rather than being a simple curve, the area receiving focused sound will be relatively small. Of course, as the talker moves, the region receiving the focused sound will move also. In addition to creating uneven distribution of sound in the audience area, this focused reflection will arrive later than the direct sound. If the difference in arrival time exceeds the Haas interval (~50 ms), an audible echo will result. This effect shows clearly in the Echo Speech map of the quasi-elliptical auditorium shown in **Photo 10**. Since Echo Speech depends upon both the strength of the each compared to the direct sound, and the delay of the echo, larger venues are more susceptible to increased Echo Speech caused by focused reflections.

I once consulted on the design of a cathedral in which the baptistry font was placed within the curve of a semicircular section of glass wall. Computer modeling showed that the sound of the priest's voice would be focused in a small area between the font and the front of the seating area. A similar design is shown in pan view in **Figure 3 [4]**. Point O is the center and focal point of the semicircular glass wall area. The voice of the priest standing at point P would have been focused at point P′. This mismatch between the priest's visual position and

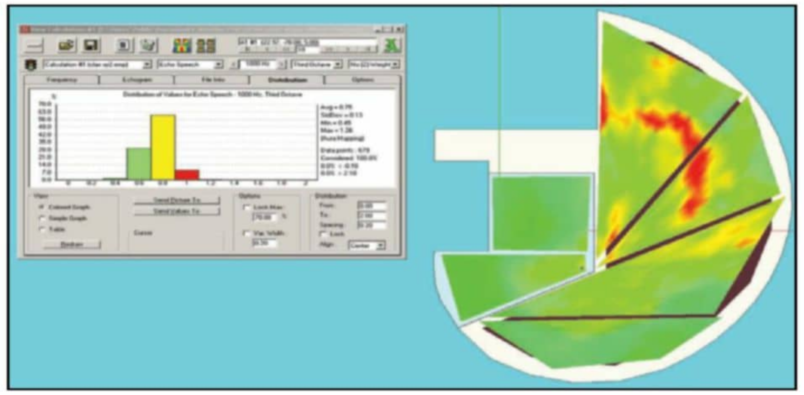

Photo 10: This map shows serious areas of echoes.

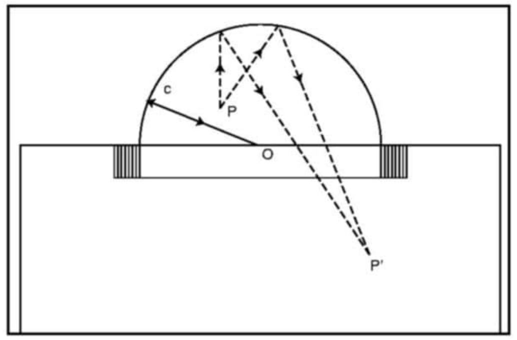

Figure 3: The curved glass feature at the top of this drawing is a sound-focusing surface.

his aural image would have been distracting for congregation members. We were able to persuade the architect to modify the design in order to prevent this problem.

When, due to owner or architect preferences, the visual effect of curved walls is required, the negative acoustical consequences can be reduced by making the wall area of plane segments rather than a smooth curve. This has the effect of introducing astigmatism into the acoustical focusing, somewhat mitigating its effect. Applying absorbent material or acoustical diffusion to the walls can also help, although neither absorption nor diffusion completely ameliorates the focusing effects of a curved surface.

It is also true that computerized acoustical modeling software is in general incapable of simulating the effects of smooth curves; such surfaces must be approximated by segmented shapes. This approximation results in the simulation understating focusing effects.

References

[1] L. Dietsch and W. Kraak, "Ein objektives Kriterium zur Erfassung von Echostörungen bei Musik- und Sprachdarbietungen," *Acustica 60*, 1986, pp. 205-216. (In German)

[2] V. O. Knudsen, *Architectural Acoustics*, Wiley, London, 1932, p. 354.

[3] Ibid.

[4] Knudsen, p. 355.

Røde NTH-100 Headphones

Making a Late, Refined Entrance

With the NTH-100, microphone specialist Røde launches its first product into the crowded studio headphone market. Brent Butterworth explores how the NTH-100 stacks up to the competition—in the studio and on the test bench.

By
Brent Butterworth

Photo 1: The Røde NTH-100 headphones are an iconic design, extremely durable and resistant, exactly as they need to be. All the important components are also user-replaceable and can be removed and swapped out in seconds. The fully adjustable headband with a unique FitLok locking system is a very distinctive feature that will be appreciated over time.

It's not surprising that Røde—a company long revered for reasonably priced microphones—decided to introduce a set of headphones. There's a long tradition of companies specializing in microphones and headphones: AKG, Audio-Technica, beyerdynamic, and Sennheiser being the best-known examples. But I assume Røde's decision had more to do with the exploding project studio market than mere brand-burnishing.

Project studios—small home studios where musicians and other content creators can work without incurring the high costs of professional studios—have been growing in popularity for the last two or three decades. But then came the explosion of two new media—YouTube and podcasts—and the growth of a new "creator class," many earning a good (sometimes great) income making video and audio programs at home or on the go. And then came the COVID-19 pandemic, which inspired many suddenly unemployed musicians to develop recording and producing skills, so they'd have something to do after losing their gigs.

Count me among them. As a jazz bassist gigging regularly with two bands and subbing for several others, I lost all my gigs when the pandemic hit, so I began working on recording projects with saxophonist Ron Cyger, whose band I play in. We posted some videos to YouTube and attracted the notice of Angela O'Neill, a Los Angeles, CA, bandleader and singer who was using her free time during the pandemic to start a record label, and invited us to make an album under the name Take2—an effort that inspired us to spend a substantial chunk of change on new studio gear.

The $149 Røde NTH-100 headphones might be a little late to the pandemic "party," but it's a sure thing that the millions of new content creators will be needing to upgrade and expand their gear as their skills (and with any luck, their revenue) increase.

The Headphones

With so many well-accepted studio headphones in the NTH-100s' price range—some dating back to the days before ProTools—creating a new set of studio cans with marketable differences isn't an easy task. But after one look at the NTH-100s, you could probably pick them out of a pile of studio headphones at 3 meters, thanks to their distinctive design.

While most over-ear headphones have a roughly oval shape, the NTH-100s are more triangular (**Photo 1**), designed loosely along the shape of the human pinna. According to Røde's press release, comfort is the main intent of this design, although the roughly triangular shape is said to eliminate standing waves and resonances. The earpads are made from memory foam with a layer of Røde's proprietary CoolTech gel, covered by Alcantara synthetic suede (**Photo 2**). Ideally, these design decisions would make it comfortable for content creators to wear the NTH-100s for hours without discomfort—a feature that's especially important for podcasters, who may indulge in a couple of hours of discussion that they edit down to 30 or 45 minutes.

Another distinctive twist is the ability to distinguish each set of headphones by color—either through tiny plastic ID rings (**Photo 3**) included with the headphones, or with optional colored cables in pink, green, orange, and blue. The included cable measures 2.4m; a 1.2m cable, suited for mobile use, is optional.

Not surprisingly, Røde's press release touts a unique driver design, with a custom voice coil surrounding a four-layer aluminum former, attached to a three-layer Mylar diaphragm (**Photo 4**). But as someone who tests 50 to 60 headphones and earphones a year—most of which claim advantages because of their driver design—I'm skeptical of any performance claims related to driver design.

Using the NTH-100s

As so much of Røde's pitch focuses on comfort, I figured I should address that first. I worried that the earpad's relatively small cavity for the ears would bother me, considering I have slightly larger than average ears. I like the triangular shape a lot, because it sort of forces the headphones into the right position over your ears, which should reduce variations in frequency response due to minor positioning differences. I did find myself wishing the ear cavities were about 10% larger, though—I could feel them pressing against my ears in places—and I do prefer the cooler feel of protein leather earpads to Alcantara. But after about 15 minutes of use, I forgot about these relatively minor differences and felt completely comfortable with the headphones. The NTH-100s had what felt to me like the perfect amount of clamping pressure, too; they didn't place uncomfortable pressure around my ears, but even shaking my head violently didn't dislodge them. I think most people would probably find this design comfortable and practical for long recording and listening sessions.

The NTH-100s aren't as efficient as some of the other studio headphones I had on hand (e.g., the AKG K371s), but they did play plenty loud enough when plugged straight into my Samsung Galaxy S10 smartphone, so they should make for an excellent companion on daily commutes.

I also liked that the NTH-100s feel ruggedly built, thanks in large part to their metal headband and solid-feeling, thick plastic enclosures—some

Photo 2: The choice of Alcantara microfiber material (with the touch of suede) for the earcups and headband cushions feel luxurious and are a great choice for long sessions, complemented with very nice-fitting CoolTech earcups that do reduce wearing fatigue and provide very good isolation.

Photo 3: It's all about the ecosystem. The NTH-100 package comes with a simple storage pouch, 1/4" screw-on adaptor, color-coded rings and one 2.4m (7'10") straight black cable. Optional color-coded cables are available to order.

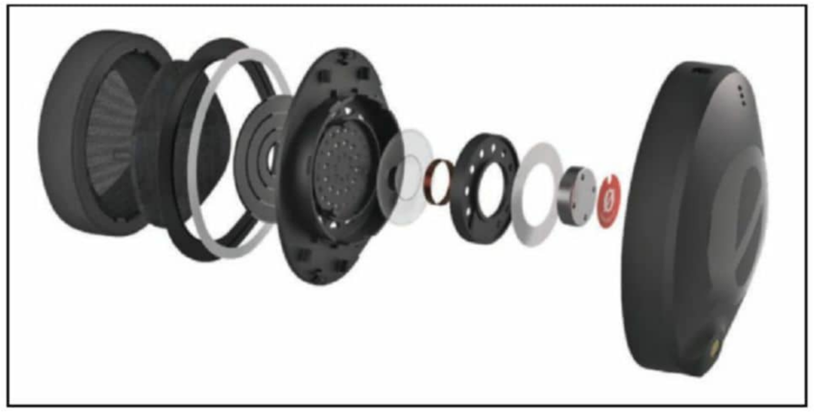

Photo 4: The NTH-100's drivers are based around a custom voice coil with a four-layer ultra-high-tension aluminum alloy as its core, which is coupled to an ultra-stiff triple-layer Mylar diaphragm and neodymium magnet.

Photo 5: Attachments on both earcups allow the cable to be positioned where its most comfortable and convenient, with a high-quality locking connector for each attachment. Left and right earcup identification is easy thanks to luminous letters, color-coding and braille. The earcups also feature high-density memory foam for excellent noise isolation.

Photo 6: Røde enters the headphone category with the NTH-100 professional headphones. The Australian company has invested a lot of time and R&D resources to enter this competitive market, and its first headphones are a refined, completely original design with very high-quality details and superior comfort. The contoured earcups are precision-engineered to be comfortable and acoustically optimized.

Resources

Røde, www.rode.com

Take2, take2.hearnow.com

competing headphones, such as the Audio-Technica ATH-M50X and KRK KNS 8402, have a plasticky feel and creak when twisted. Of course, the way the headphones feel doesn't necessarily indicate their longevity, but the NTH-100s definitely don't feel cheap—I'd feel good about handing them to other musicians I've working with during a recording session.

The cable is a proprietary design, with a standard 3.5mm stereo plug on the source end and a twist-lock 3.5mm on the headphone end. This means the cable can't be accidentally pulled loose, but you'll probably have to purchase spares directly from Røde. Another nice touch is that the cable can be plugged into the left or right earpieces—something that's very convenient in certain situations (**Photo 5**). Right-handed guitar and bass players will probably like it because the cable can hang off to their right side rather than to the left (as would be the case with most studio headphones), and will be less likely to interfere with their instruments.

The Sound

My attitude about studio headphones has always been that none of them are perfect—but the best of them, despite their differences and minor idiosyncrasies, all work well once you get used to their sound. I'd definitely rank the NTH-100s among the "best of them."

To see how the NTH-100s compare with my usual studio headphones, I compared them with the AKG K371s, which are noted for their close adherence to the Harman curve (which conveniently makes them sound very much like my JBL 305P MkII monitor speakers, also made by Harman). To do this, I opened up the Reaper digital audio workstation (DAW) to play some of the mixes from the Take2 album to see if I'd have made any different mixing decisions. While the two headphones sounded different, I didn't get the impression that I'd have made substantially different decisions.

The NTH-100s definitely have more bass than the K371s, and the bass also sounds a little fatter and less defined through the NTH-100s—so I'd probably have reduced the bass level and also boosted the bass tracks around 1.5kHz to give them a little more "growl." The NTH-100s also sounded a little softer in the region around 3kHz to 5kHz than the K371s did, so I'd probably have boosted that range by a decibel or two to bring out the saxophone, flute, electric guitar, congas and bongos.

So I'd guess that if I'd have recorded Take2 using the NTH-100s—going back and forth between the headphones and the JBL monitors—the mix might

have been a little less bassy, and would probably have had more zip in the lower treble. But I suspect that the mastering engineer, Allan Tucker of Foothill Digital, might have negated those decisions when he did his EQ. Either way, I'm confident that the end result would have been similar, and I'd have been happy with the final product.

I also compared the NTH-100s with the K371s, the KRK KNS 8402s and the Audio-Technica ATH-M50Xes, using recordings I have used for many years in headphone and speaker reviews. To my ears, the NTH-100s and the K371s felt like "different flavors of the same thing"—I enjoyed listening to both, and while the NTH-100s had more prominent bass and softer lower treble, my ears were easily able to adjust to the difference.

This wasn't the case with the ATH-M50Xes—which still sounded good (they're a very popular model with audio pros, many of whom produce great-sounding recordings), but had a more prominent, insistent bass that actually made the bass notes seem to linger longer. They also seemed to have a slightly recessed midrange, which made voices fade deeper into the mix, and they also apparently added a bit of extra emphasis in the mid-treble that made voices sound a little edgier. The KNS 8402s seemed to ramp up on the ATH-M50Xes' sonic formula, with more bass and treble, and less midrange energy.

I enjoyed listening to music with all of these headphones. But if I had to do a mix on any of them, I'd pick the K371s or the NTH-100s. Of course, the general rule in audio production is that it's wiser to mix on speakers rather than headphones, but in an era when content creators are cranking out such a colossal amount of content—sometimes from hotel rooms, airplane cabins, or who-knows-where—it's important to be able to mix on headphones in situations where you don't have the luxury of a good set of studio monitors (**Photo 6**).

Measurements

I performed measurements using a GRAS Model 43AG ear-cheek simulator, an RA0402 high-resolution ear simulator/coupler, an Audiomatica Clio 12 QC audio analyzer, and an M-Audio Mobile Pre USB interface with TrueAudio TrueRTA real-time analyzer software.

I started with the most basic measurement for headphones: frequency response. **Figure 1** shows the response in the left and

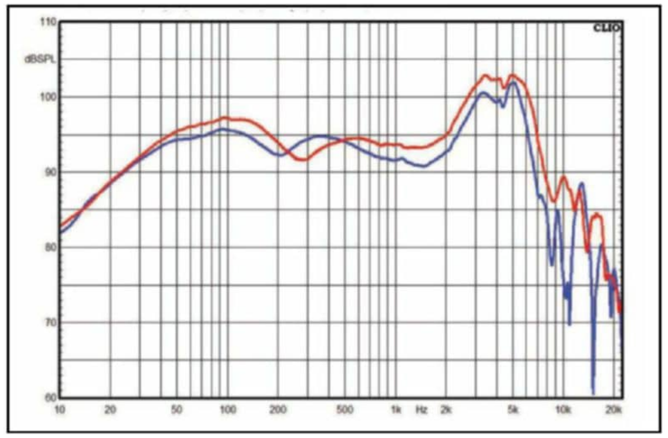

Figure 1: Frequency response of the NTH-100 in left (blue) and right (red) channels, referenced to 94dB at 500Hz

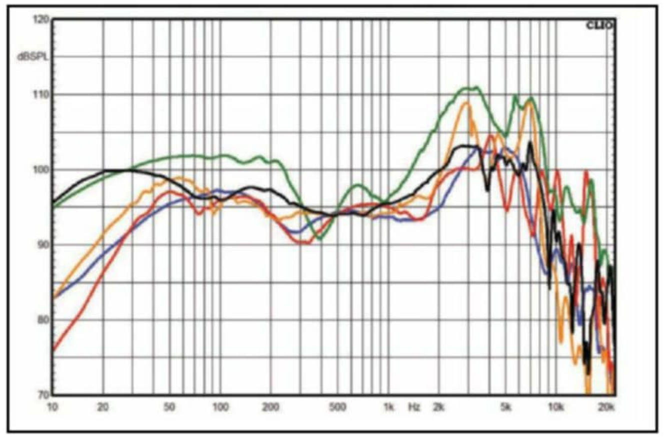

Figure 2: Frequency response of the NTH-100 (blue), AKG K371 (black), Audio-Technica ATH-M50X (red), Sony MDR-7506 (orange) and KRK KNS 8402 (green)

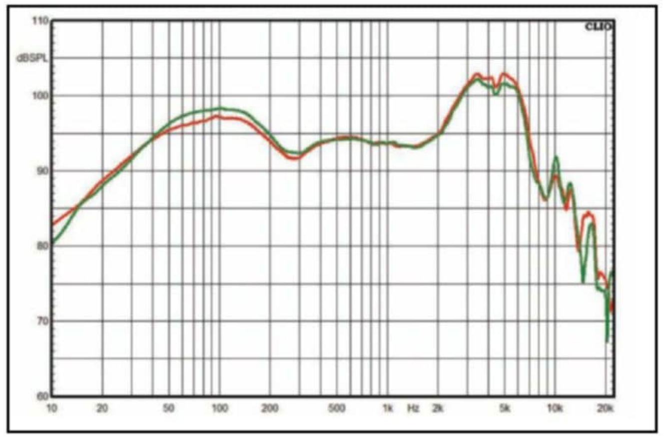

Figure 3: Frequency response of the NTH-100 with 5Ω (red) and 75Ω (green) sources

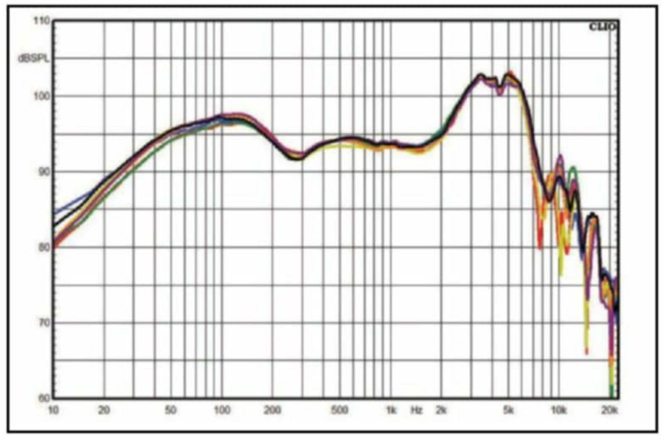

Figure 4: Frequency response of the NTH-100 with multiple reseatings on the ear-cheek simulator

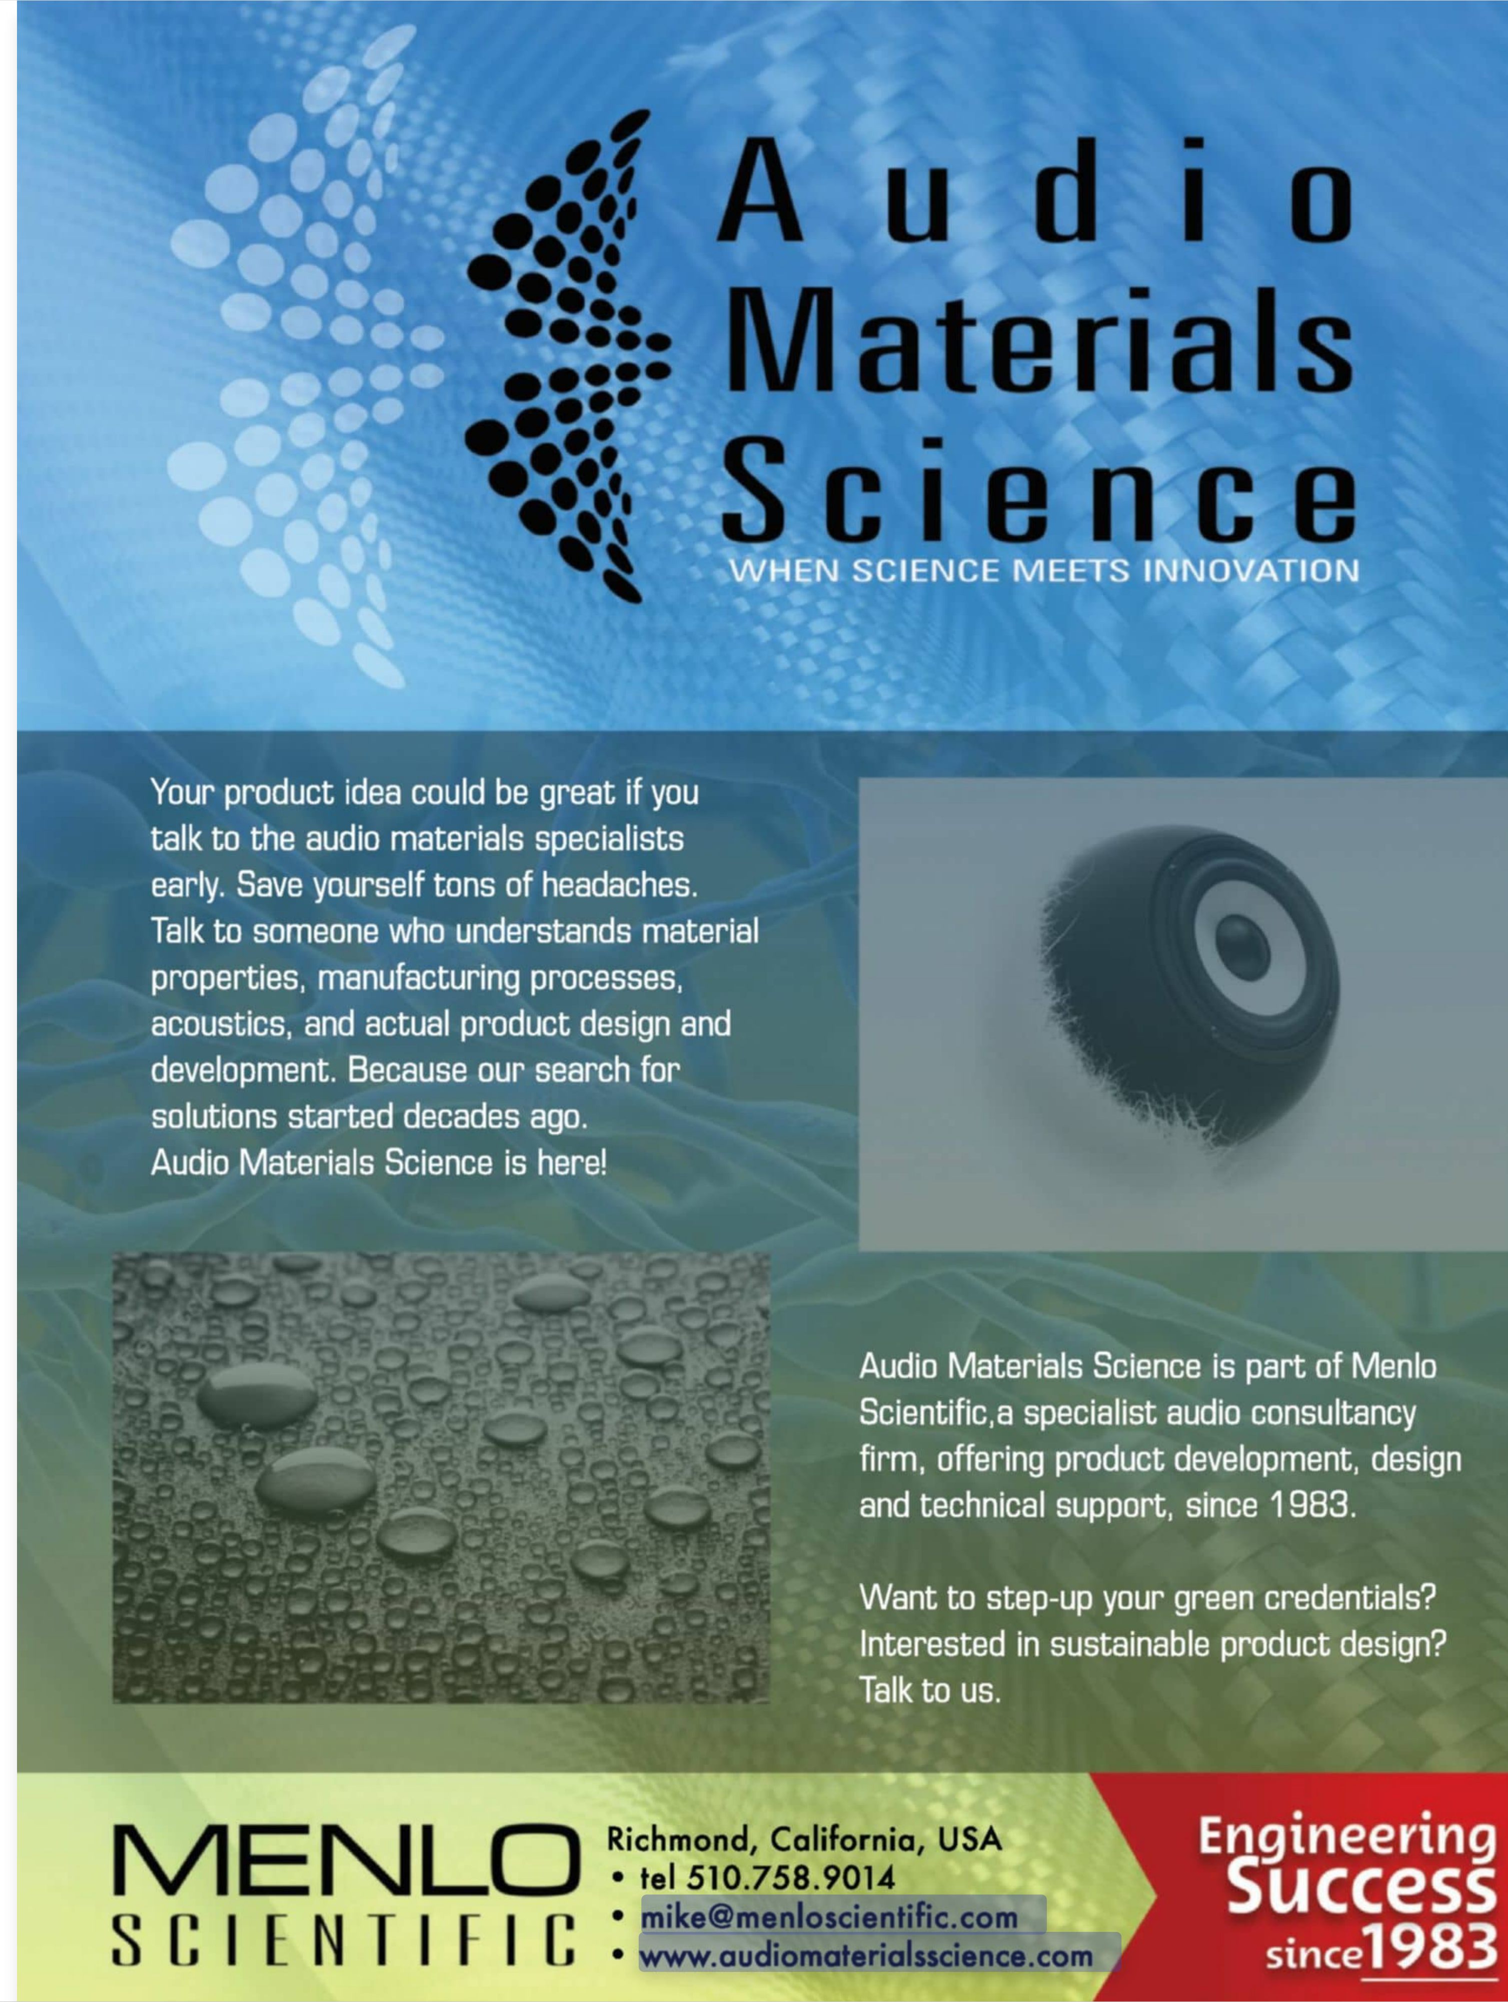

right channels. This is a typical result for headphones of this sort. I also made the same measurement with a 1mW signal calculated for the NTH-100s' 32Ω rated impedance, between 300Hz and 3kHz, which showed an average sensitivity of 102.7dB—not super-high, but high enough that no source device should have trouble driving the NTH-100s to high volumes.

Figure 2 is more revealing. It's a comparison of the frequency response of several professional studio headphones that compete with the NTH-100s. The NTH-100s appear to be similar here to the AKG K371s and the Audio-Technica ATH-M50Xes.

Figure 3 shows how the frequency response is affected when the source impedance is switched from 5Ω (the output impedance of the Musical Fidelity V-CAN amplifier I used) to 75Ω (with 70Ω of resistance added in series with the V-CAN), which is sometimes encountered in professional headphone amps. It's a small difference—a boost of about 1dB in the bass and a cut of about 1dB in the treble—which suggests that any differences in tonal balance when switching from one amp to another will be small.

The efficacy of the NTH-100s' ear-hugging design can be seen in **Figure 4**, which shows the difference in response when the headphones are moved about 6mm up, down, forward, and backward on the ear-cheek simulator, then removed and reseated several times. There's very little difference, which suggests that the NTH-100s' sound will not vary significantly depending on fit and positioning.

Figure 5 shows the spectral decay—or "waterfall"—of the NTH-100s, which reveals resonances. This is a very clean-looking plot, with only some mild resonance in the bass that's damped within a couple of cycles.

Figure 6 shows the total harmonic distortion (THD) of the NTH-100s at 90dBA and 100dBA. **Figure 7** shows how the result at 100dBA compares with the distortion of the AKG K371s, Audio-Technica ATH-M50Xes and KRK KNS 8402s. Clearly, the distortion of the NTH-100s is a little higher than the competition, but still, it's only 3% between 100Hz and 200Hz, and 5% below 40Hz. Considering the relatively low audibility of headphone distortion compared with amplifier distortion, the low frequencies in question,

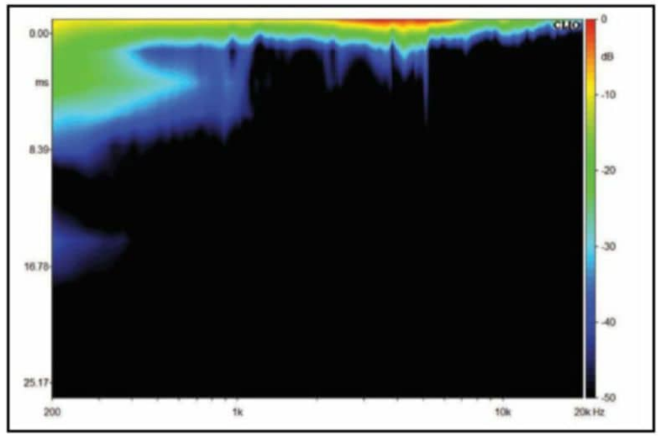

Figure 5: Spectral decay of the NTH-100

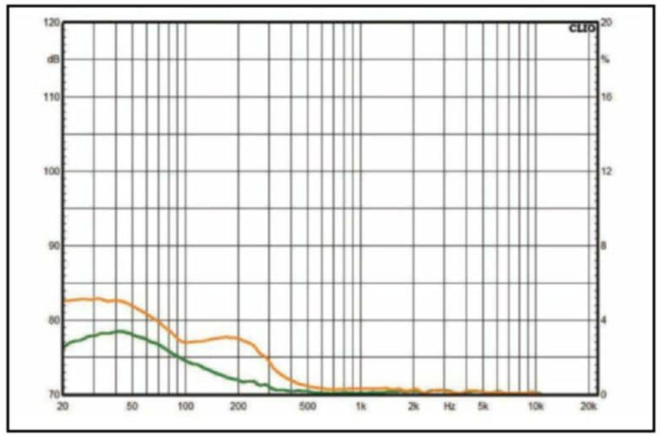

Figure 6: Total harmonic distortion vs. frequency of the NTH-100. Playback levels set at 90dBA and 100dBA using pink noise.

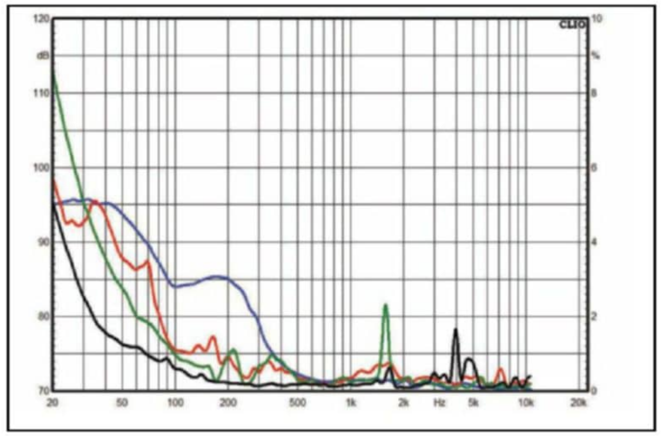

Figure 7: Total harmonic distortion vs. frequency of the NTH-100 (blue), AKG K371 (black), Audio-Technica ATH-M50X (red) and KRK KNS 8402 (green)

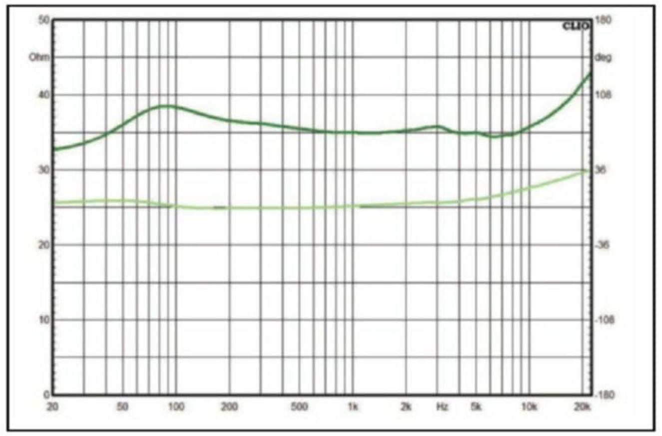

Figure 8: Electrical impedance magnitude (dark green) and phase (light green) of the NTH-100s

and the extremely high playback levels involved, it's very unlikely this distortion would be audible.

Figure 8 shows the impedance magnitude and phase of the NTH-100s. It's mostly flat, and generally hovers around 35Ω.

Figure 9 shows the ability of the NTH-100s to isolate the listener's ears from outside sounds. The measurement was taken by playing uncorrelated pink noise at 85dB SPL from a six-speaker system. The baseline level is shown by the red line; the other lines show how much the headphones attenuate this noise. The isolation provided by the NTH-100s appears to be adequate but unexceptional.

Conclusion

While Røde's promotional literature bills the NTH-100s as the "perfect headphones," I wouldn't say they're perfect—but I wouldn't say any audio product is perfect. I would say, though that the NTH-100s are at least the equal of any competitor in their price range, and an outstanding choice for YouTubers, podcasters, musicians, studio pros, and audiophiles.

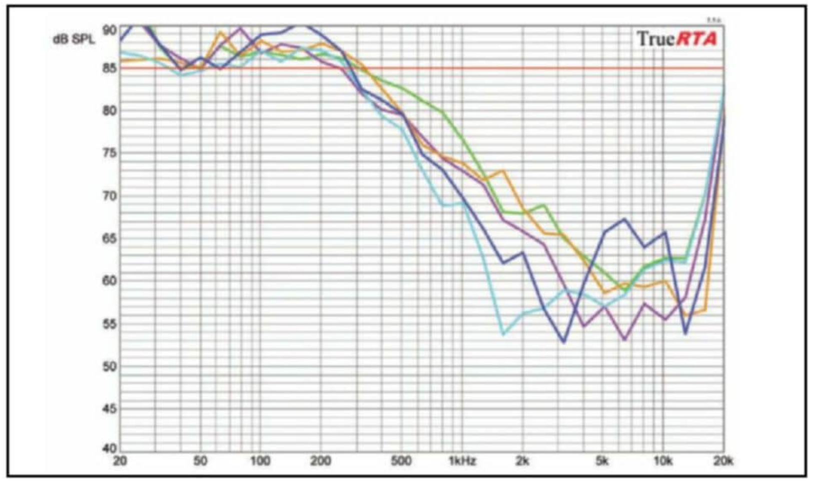

Figure 9. Isolation of the NTH-100s (green trace) compared with the AKG K371s (purple), Sony MDR-7506 (orange), KRK KNS 8402 (cyan) and Direct Sound Serenity II (blue).

About the Author

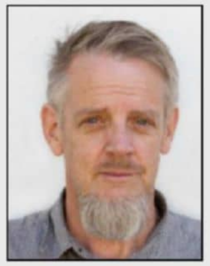

Brent Butterworth has been a professional audio journalist since 1989, and has measured and evaluated thousands of audio products. His headphone measurements and reviews currently appear on soundstage.com and wirecutter.com. He has served as editor-in-chief of *Home Theater* and *Home Entertainment* magazines, contributing technical editor for *Sound & Vision* magazine, senior editor of *Video* magazine, and reviews editor of *Windows Sources* magazine. He has also worked as marketing director for Dolby Laboratories, and as a consultant on the design and tuning of speakers, subwoofers, and headphones.

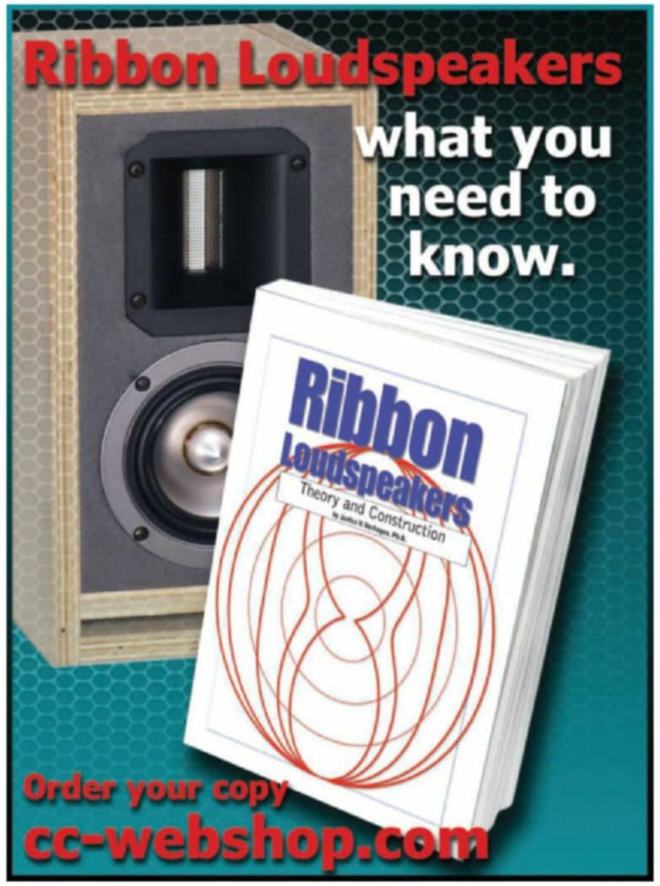

How Signal Processing and Filter Bank Design Impacts Sound Quality in Hearing Aids

This article discusses the nuances of signal processing and filter bank design, exploring the most impactful decisions in filter bank types specific to hearing aids, and how ultimately each approach helps dealing with the hearing challenges that people face.

By
Dana Helmink
(Senior Director, Audiology/Clinical Education, Widex USA)

Modern hearing aids are marvels of digital audio processing. They take analog sound waves, convert them to digital signals, process them in various ways to overcome a wearer's particular form of hearing loss, then convert them back into sound waves played through the hearing aid's earpiece. Of course, digital audio processing is nothing new, but when it's performed in hearing aids, it's meant to have a demonstrable impact on wearers' lives. Therefore, doing it well is important.

By 2050, the World Health Organization estimates nearly one in four people globally will experience some form of hearing loss. About 700 million of them will need to take corrective action. Among other risks, research indicates that hearing loss can lead to social isolation, and studies from around the world show that social isolation can damage cognitive function. But it's important to understand that for many, the challenge isn't just about hearing more (volume); it's about hearing better (clarity). It's about picking up the voices of friends and family in a noisy restaurant, for example, or understanding clearly and naturally the words people say. And sometimes, hearing aid technology actually hinders that clarity.

That's because signal processing in hearing aids takes time—about 5 to 8 milliseconds. This delay in processed audio, when coupled with direct, immediate sound that comes through a hearing aid's venting, degrades sound quality for the wearer. Known as the comb filter effect, this mixing of processed and direct sound in the ear canal can cause distortion, "tinniness," and general discomfort when listening over time.

This is where advances in signal processing come into play—and where the choice of signal processing technique can positively impact the natural sound experience for hearing aid wearers. The goal is to minimize delay and quash the comb filter effect, resulting in more natural sound and better hearing. Through a combination of high sampling rate and time-domain filter banks, this natural sound in hearing aids is achievable.

Nuances of Signal Processing in Hearing Aids

Overall, the processing pathway in digital hearing aids is similar across solutions. A microphone picks up the sound and converts it from analog to digital. An analysis filter bank splits the signal into multiple frequency bands, or channels, readying them for further digital signal processing, such as compression and noise reduction. After processing, the signal is recombined by a synthesis filter bank and played out through the receiver in the hearing aid.

This common approach to hearing aid design illustrates the important role of filter banks in digital signal processing. Choosing the right filter bank, which we will discuss in a moment, is key. Also important is the role of sampling, the process of converting the continuous analog sound signal into discrete parts (samples) for processing. As is well known, CDs and many music streaming services use a sampling rate of 44.1kHz (44,100 samples per second). For hearing aids, audio sampling rates typically range between 20kHz and 33.1kHz.

Of course, sampling rate is related to the frequency that can be represented in the signal, so

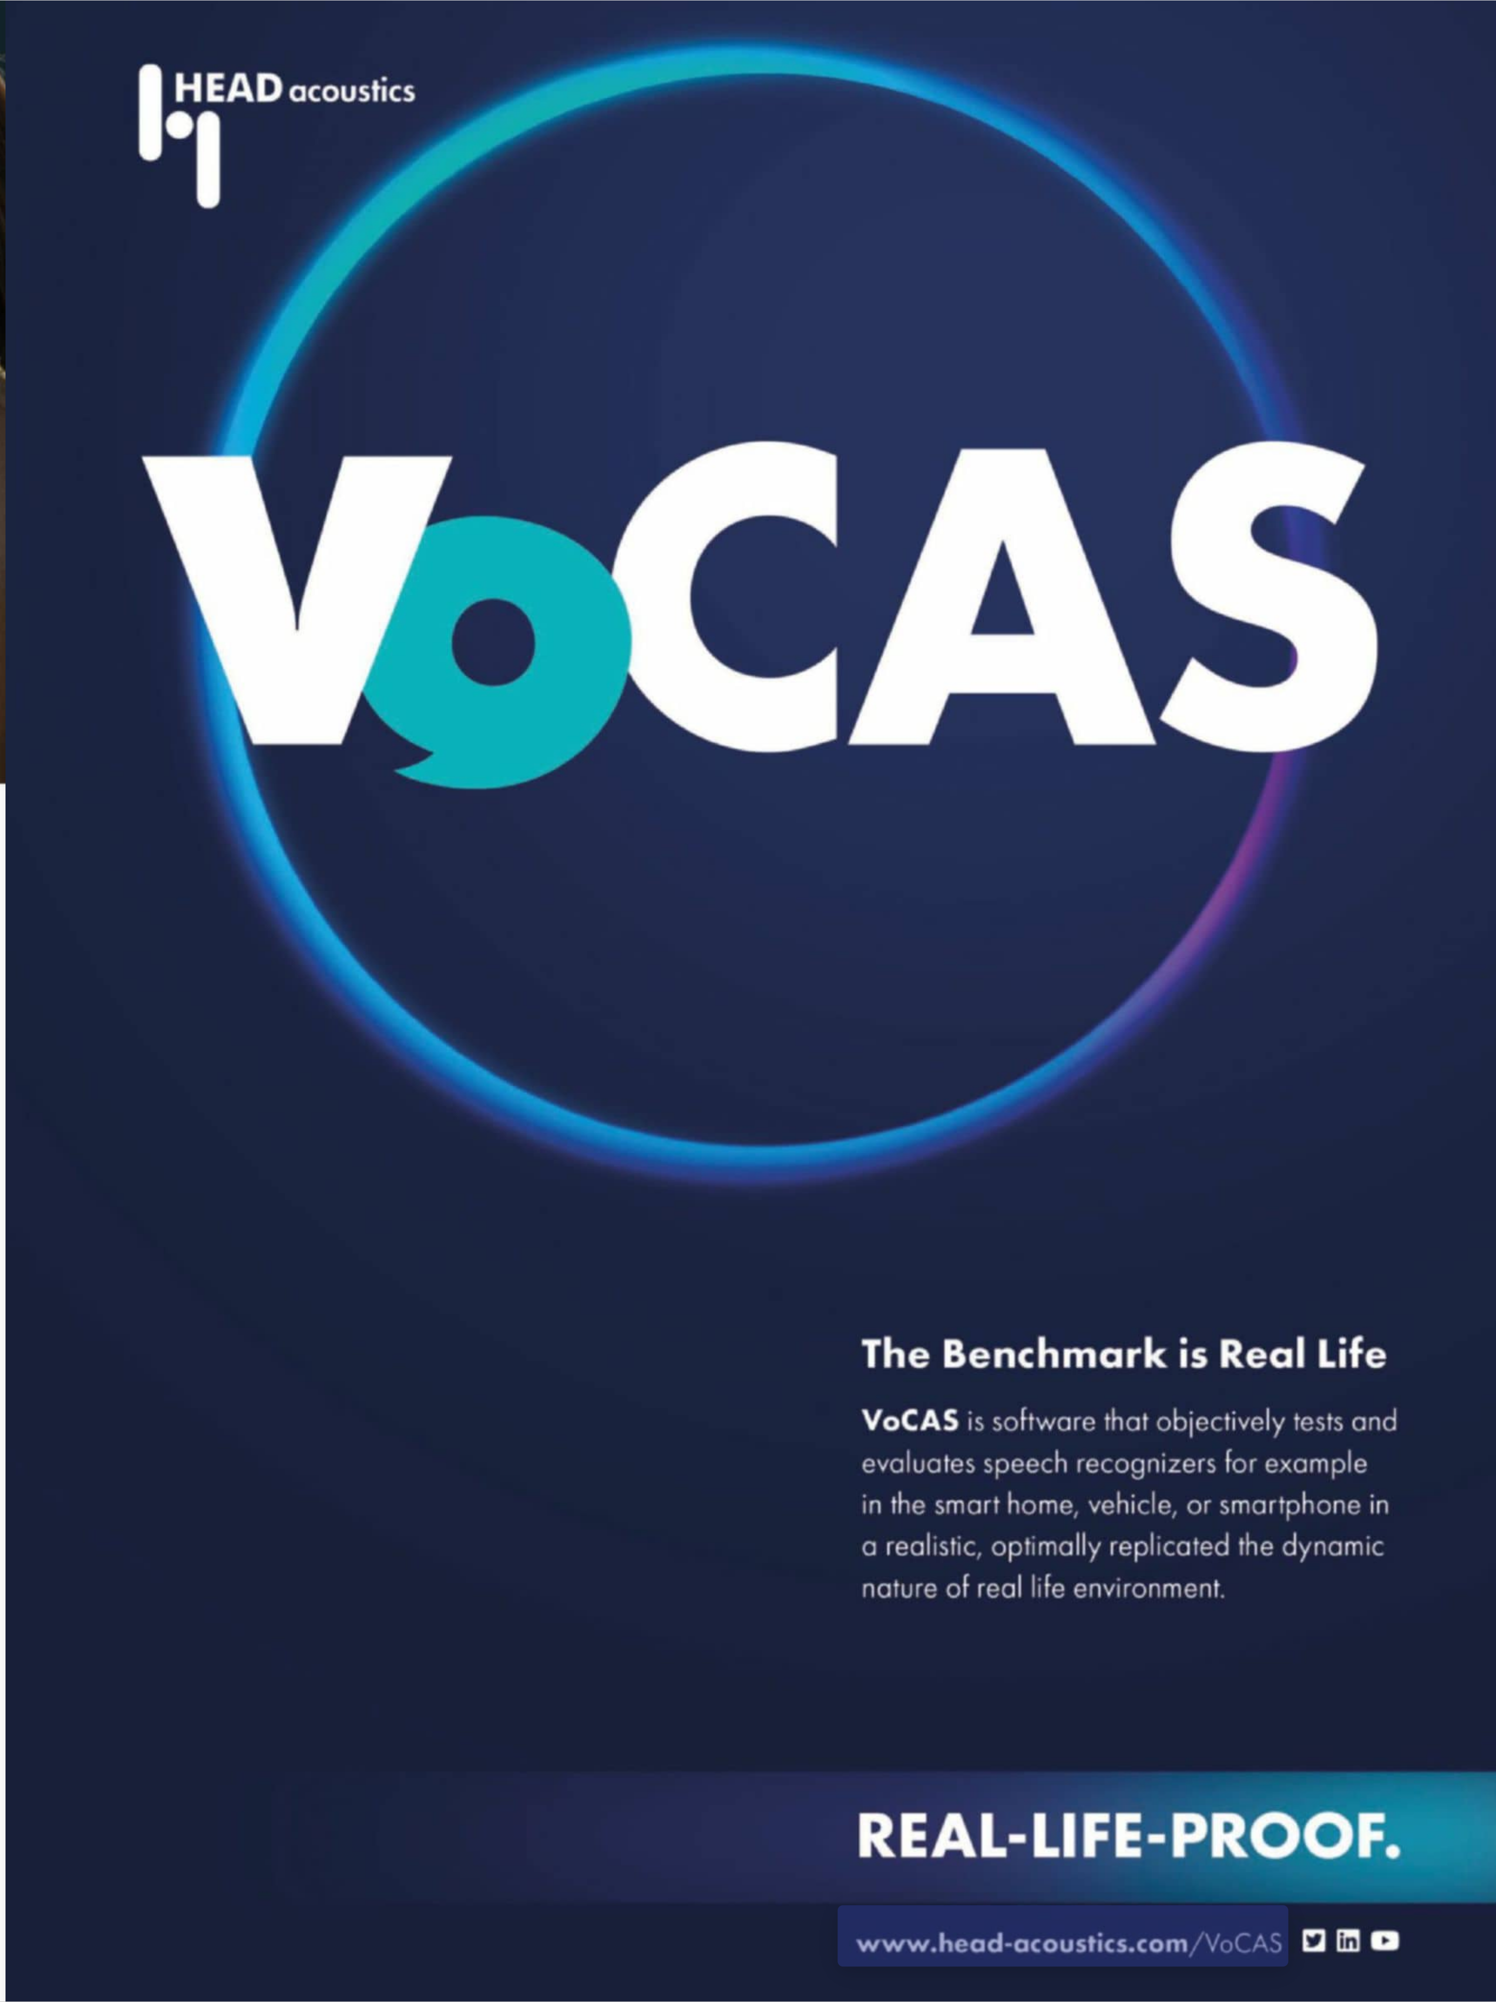

that the highest frequency that can be represented is half the sampling rate (the Nyquist frequency). This means that in hearing aids, a 20kHz sampling rate has a Nyquist frequency of 10kHz, allowing frequencies up to 10kHz, while a 33kHz sampling rate allows frequencies up to 16.5kHz.

These are sampling rates at the input stage; later in the processing pathway, the signal may be down-sampled to save processing capacity, before being up-sampled again. Although other factors, such as receiver roll-off, may limit the frequency range of the output, higher fidelity in the high frequencies is essential for addressing certain forms of hearing loss. Therefore, a higher sampling rate has the best chance of creating that natural sound experience. It can also offer other sound quality benefits, including lower processing delay and more gently sloping filters.

Choosing a Filter Bank

As mentioned earlier, perhaps the most impactful design decision in hearing aids is which type of filter bank to use. Frequency-domain filter banks are popular in hearing aids, often for good reason. But when it comes to creating more natural sound, reducing processing delay, and delivering better hearing, time-domain filter banks are the superior choice.

The difference between the two has to do with a trade-off in processing between time and frequency resolution. In any filter bank, the time resolution of each frequency band is proportional to the width of that frequency band. Broadly speaking, this means that a bandwidth that is five times wider has a time resolution that is five times higher. Conversely, a narrower bandwidth means poorer time resolution.

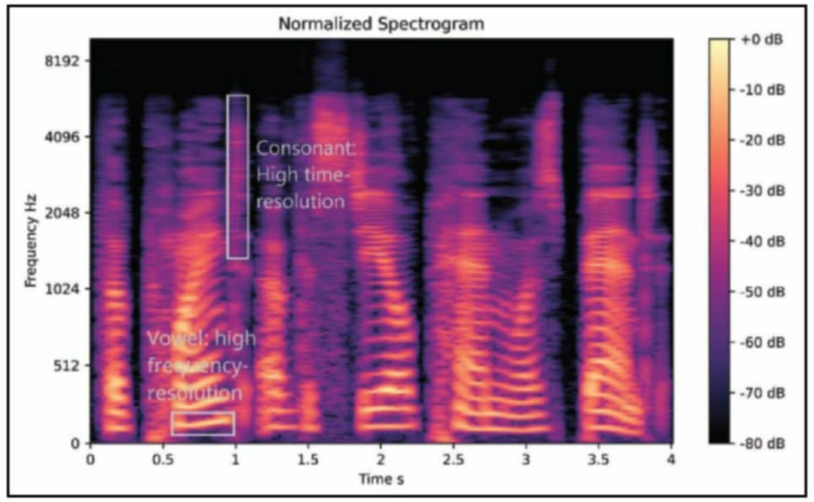

Figure 1: Speech spectrogram illustrating a vowel (lower box) that requires higher frequency resolution and a consonant (upper box) that requires higher time resolution

Frequency-domain filter banks limit the system to frequency bands of the same width across the entire frequency range. In practice, this means that all bands are relatively narrow because bandwidth is set based on the bandwidth needed for the lowest frequencies, where the ear's frequency sensitivity is the highest, and the bands therefore must be narrow. Because of the trade-off between time and frequency, this also means that all bands—both high- and low-frequency—operate with the same, relatively poor, time resolution.

On the other hand, time-domain filter banks offer the flexibility to use filters that may vary in bandwidth. This means hearing aid designers can set bandwidths any way they like. With a frequency-domain filter bank, all bands have the same width and time resolution, but with a time-domain filter bank, the bands can vary in width and, therefore, also in time resolution. This is important when it comes to a hearing aid use case.

Mimicking How People Hear

What does it mean in practice? With a time-domain filter bank, hearing aid designers can make a filter bank with narrower bands at the lower frequencies and broader bands at the higher frequencies, thus keeping the same trade-off between time and frequency in the filter bank as humans have in their auditory systems. This filter bank design mimics the logarithmic way the ear works, with higher frequency sensitivity at lower frequencies and lower sensitivity at higher frequencies.

Admittedly, frequency-domain filter banks can ostensibly mimic the logarithmic way the ear works by collapsing bands at the higher frequencies and reducing the number of bands overall. But that doesn't address challenges of time resolution, which can impact speech intelligibility.

When we look at a speech spectrogram (**Figure 1**), we can visualize the challenge. Spoken consonants are short in duration, but high and wide in frequency. Spoken vowels are longer in duration but lower and narrower in frequency. This means the signal would be best represented by filters where the high frequencies have high time resolution while the low frequencies have better frequency resolution—something better achieved using time-domain filter banks.

Avoiding Down-Sampling for Better Sound Quality

One more thing to consider about filter bank design in hearing aids is the role of down-sampling. For wearers, hearing aids are vital lifestyle devices. Ideally, they last a long time before needing to be recharged, which means power-consuming functions like signal processing need to be optimized. Down-sampling can help accomplish this, but it can also introduce annoying noise artifacts.

Frequency-domain filter banks with narrow filters make down-sampling straightforward, and fewer samples mean less power consumption. With a time-domain filter bank, the benefits of down-sampling are minor and don't outweigh the risk of artifacts that are inherent to non-linear processes like down- and up-sampling. Yes, a time-domain filter bank means we need other available solutions for keeping power consumption sufficiently low, but by avoiding down-sampling and preserving the original input signal, we achieve higher fidelity, reduce the risk of artifacts, and improve natural sound quality.

Indeed, choosing a time-domain filter bank can make it more challenging to optimize power consumption and certain other aspects of signal processing—though both are very achievable. But ultimately, a time-domain filter bank better complements the human auditory system and helps optimize sound quality. Combined with a high sampling rate, which significantly reduces artifacts in the signal that would otherwise need to be fixed, processing delay can be driven down close to about 2.5 milliseconds, virtually eliminating its effect on hearing.

In other words, when it comes to signal processing in hearing aids, choosing frequency-domain filter banks may be best for the hearing aid manufacturer. But a time-domain filter bank is best for the person the hearing aid is designed to help.

About the Author

Dana Helmink is Senior Director, Audiology/Clinical Education at Widex USA. She earned her Master of Arts in Audiology from Northern Illinois University and her Doctor of Audiology degree from Salus University. She previously worked in clinical audiology in Illinois, where she also served more than 10 years on the Illinois Academy of Audiology board of directors. With more than 15 years in product management, Dana applies her experience in user-centered design and usability testing to develop innovative training programs.

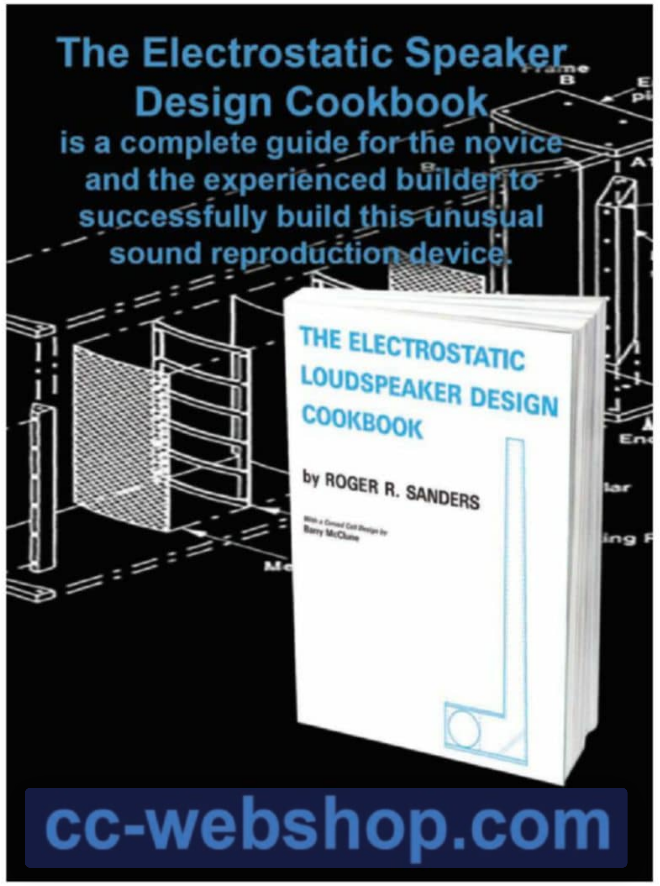

ANC Headset Testing

By
H. W. Gierlich
(HEAD acoustics GmbH)

This article discusses requirements for effective performance testing of communication headsets that use active noise cancellation (ANC) technologies. It addresses test setups, testing equipment and setup conditions needed to correctly assess ANC performance.

Headsets are used with increased frequency for phone calls and multimedia applications, and the trend of working from home further accelerates the use of headsets at work. Audio/video group applications, such as MS Teams, Gotomeeting, Zoom, and others work best when used with high-quality headsets, and, especially in mobile scenarios, adaptive noise cancellation (ANC) is a prevalent and highly desired feature. Headset makers use various ANC certification schemes from several application providers to qualify their products. Implementing new technologies in headsets requires increasingly sophisticated testing and optimization strategies, starting with using the right artificial ear and requiring new, more advanced testing technologies. All this aims to simulate the average human as precisely as possible and provide the most realistic test environment to address communication conditions and typical human behavior in the lab. Many aspects must be covered; This article focusses on performance evaluation in receiving.

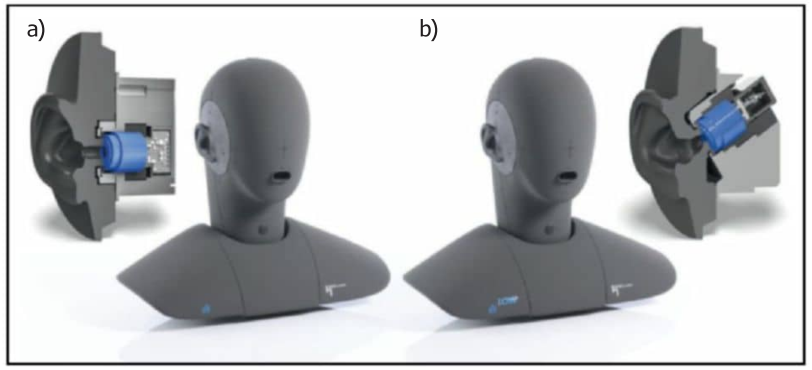

Figure 1: Traditional HMS II.3 equipped with a Type 3.3 artificial ear (a) and HMS II.3 LN-HEC equipped with a Type 4.4 artificial ear (b) are shown.

The Right Choice of Artificial Ear

A variety of artificial ears have been standardized over the last decades [1], [2]. The most important question is, as always: How realistic can measurements be, and how reproducible can one make measurements? The more simplistic an ear simulator, the more reproducible results can be expected. However, modern communication devices, especially in-ear devices, require a more realistic reproduction of the human pinna for measurements. The measured performance of such devices depends on realistic simulations of the ear canal, the ear impedance, and the mechanical properties of the complex structure and geometry of the human ear. Furthermore, especially with noise-canceling headsets, very low self-noise of the measurement equipment is required. The limitation is typically the microphone noise of the artificial impedance simulator—besides noise that may be present in the test room.

ITU-T has addressed these topics in its latest revision of Recommendation P.57. Two new types of artificial ears are described (Type 4.3 and Type 4.4). Both types model the concha and the ear canal entrance of human ears much more precisely than the former Type 3.3 and Type 3.4 artificial ears. In addition, Recommendation ITU-T P.57 describes the requirements for low noise ear simulators. The idle noise level must be below the diffuse-field hearing threshold defined in ISO 389-7 [3].

The HMS II.3 LN- HEC [12] (**Figure 1**) is equipped with a Type 4.4 artificial ear that fulfills both requirements ideally. This artificial ear conforms to the requirements for a Type 4.4 artificial ear and the low noise floor (**Figure 2**) requirements from Recommendation ITU-T P.57.

When testing headsets, especially intra-concha type headsets, using the newly defined pinna simulator is advantageous. The anatomical-shaped pinna simulator allows a most natural fit for these types of headsets since the shape of the concha bottom and the orientation of the ear canal are most human-like. When using artificial ears such as ITU-T 3.3 or 3.4 artificial ears, the 90-degree orientation of the ear canal prevents a proper fitting of intra-concha type headsets. These types of headsets often can only be positioned by bending the rubber-made ear canal of the artificial ear, thus leading to an unnatural fit.

ANC Performance

The most important measurement when determining ANC performance is the insertion loss measurement. Here, we must differentiate between passive and active insertion loss.

The passive insertion loss is measured as the attenuation of a headset when applied without ANC compared to the sound pressure level measured with no headset present.

The total insertion loss is measured as the attenuation of a headset when applied with ANC compared to the sound pressure level measured with no headset present.

The active insertion loss is the difference between both and denotes the additional loss provided by the ANC algorithm.

The setup for these measurements is shown in **Figure 3**. A Head and Torso Simulator (HATS) is used for realistic testing. Only one HATS (HATS 1 - HMS II.3 LN HEC) is used to test ANC performance alone. Most important is a realistic background noise simulation to reproduce actual noises in a laboratory environment. A procedure for recording and reproducing sound fields almost physically correct is described in ETSI TS 103 224 [4]. In this standard, various realistic background noises are included, which can be used for headset testing. Due to the nonlinear and time-variant behavior of

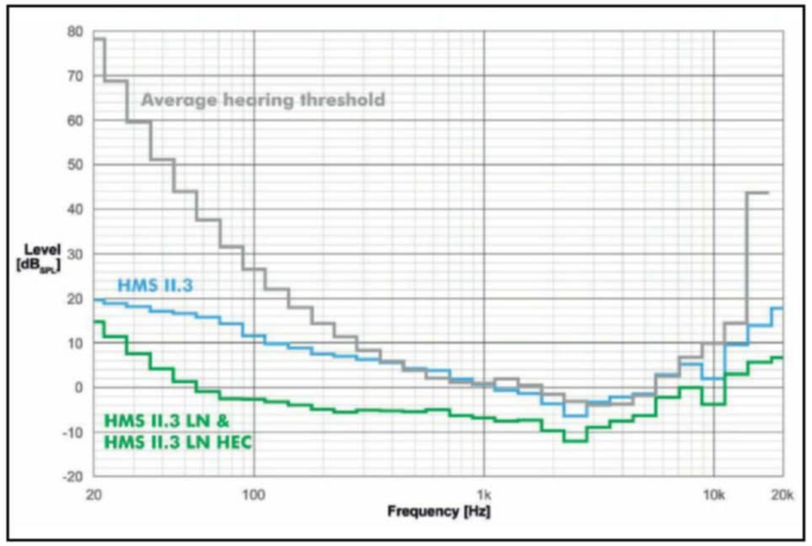

Figure 2: This graph shows the noise floor of HMS II.3 LN HEC (ITU-T Type 4.4 artificial ear) and traditional HMS II.3 (ITU-Type 3.3/3.4 artificial ear) compared to the average hearing threshold in diffuse-field conditions.

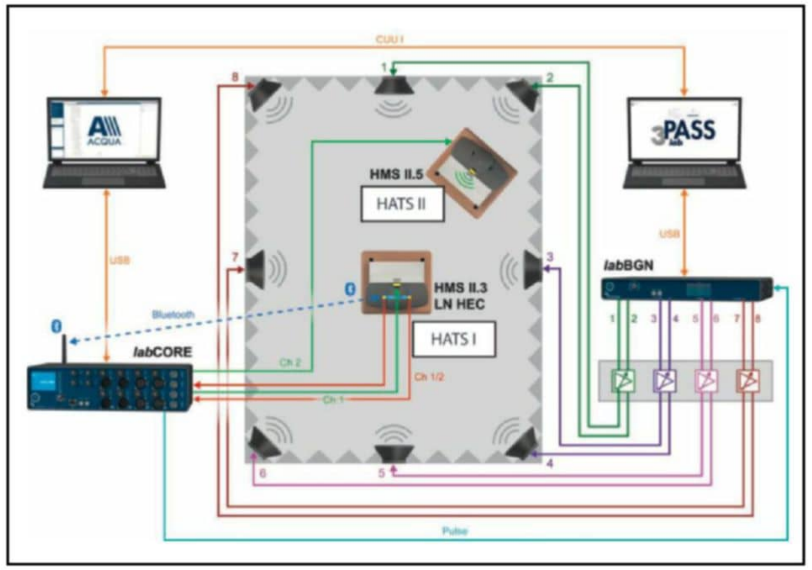

Figure 3: Here is a measurement setup for ANC headset testing. Analysis system (ACQUA, labCORE & HATS HMS II.3 LN HEC) in combination with background noise simulation (3PASSlab and labBGN).

ANC headsets, it is vital to precisely synchronize the test system with the background noise playback and conduct a conditioning sequence before starting the tests.

When measuring different headsets, positioning is crucial. To ensure the correct positioning of insert type headsets, ANC is activated and the attenuated spectrum of the background noise is monitored. Special care must be taken to achieve almost symmetrical attenuation between the left and right ear, which indicates correct positioning. Care must also be taken not to overfit by applying high pressure to the headsets, not reflecting the natural coupling to the human ears. The tests should be repeated at least five times to eliminate positioning inaccuracies, which are unavoidable and always lead to measurement uncertainties at higher frequencies. This procedure is to be applied to all types of headsets.

Some typical measurement results when conducting such tests are shown in **Figures 4-6**. Besides total insertion loss, active and passive insertion losses are indicated. Passive insertion loss for this intra-concha device is only marginal under 200Hz. This can be improved in the frequency range from 30Hz to 1kHz by enabling the ANC functionality of this headset. However, enabling ANC may decrease the performance at higher frequencies (active insertion loss between 3kHz and 10kHz). The resulting total insertion loss is shown in Figure 6, which is the sum of active and passive insertion losses.

Another effect when using a different type of background noise is shown in **Figure 7**. The resulting idle noise of the artificial ear may have a quite some influence on the measured result if the remaining noise of the measured headset is very low. In Figure 7, the measured results are compared between an ITU-T 3.3 artificial ear and a low-noise Type 4.4 artificial ear. The remaining noise floor of the standard 3.3 artificial ear leads to a virtually higher noise floor compared to the measurement using the low-noise artificial ear. The lower the self noise of the headset the more dominant this effect will get. Therefore, all types of ANC headset measurements should use an artificial ear that fulfills the low noise requirements as defined in ITU-T Recommendation P.57: 06/2021, Artificial ears [1] and conforms to the impedance requirements specified for ITU-T Type 4.3 and Type 4.4 artificial ears.

Talk-Through Performance

Many headsets provide settings that enable users to change the ANC setting to better listen to the environment (e.g., for listening to announcement systems or talkers in the room). This

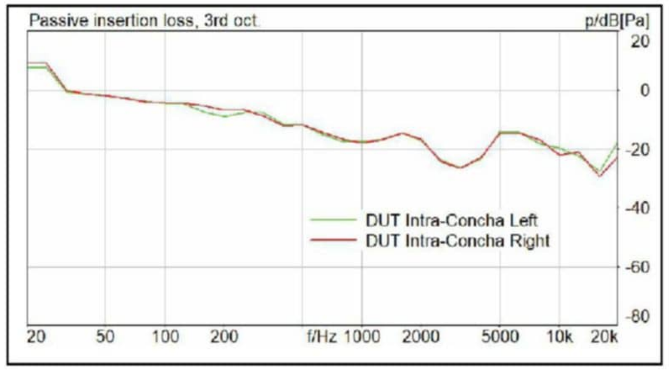

Figure 4: Passive insertion loss for this intra-concha device is only marginal under 200Hz.

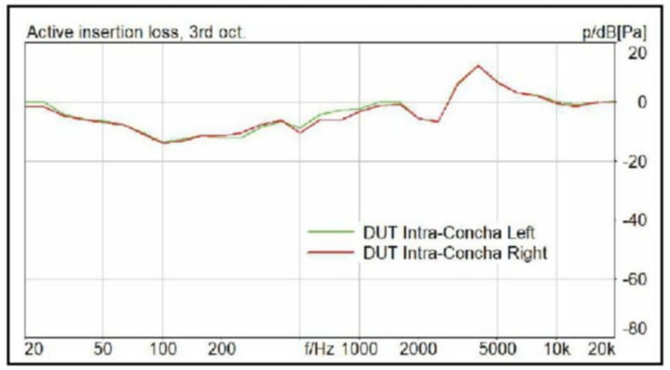

Figure 5: Enabling ANC may decrease the performance at higher frequencies (active insertion loss between 3kHz and 10kHz).

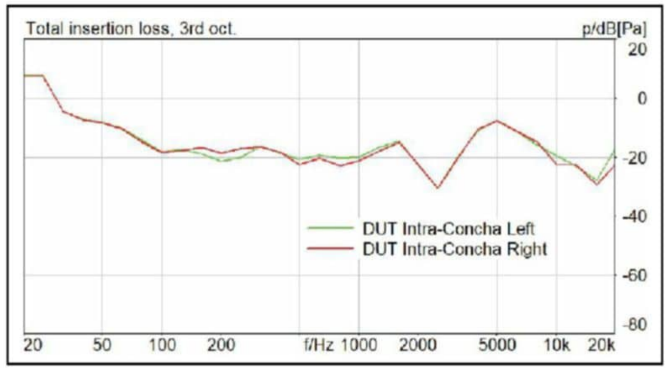

Figure 6: The resulting total insertion loss is the sum of active and passive insertion losses, measured in a diffuse pink noise sound field [4] for an intra-concha headset.

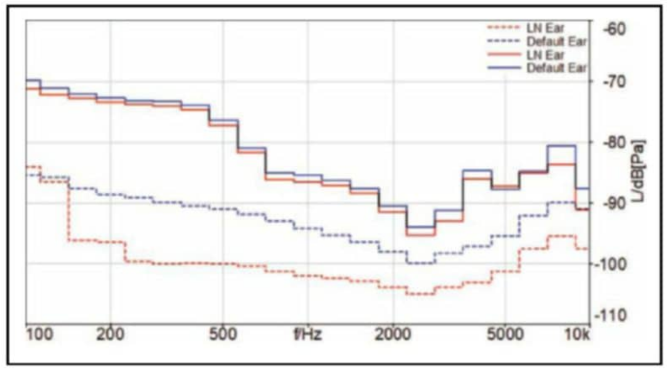

Figure 7: Self noise of an intra-concha headset is shown using an ITU-T Type 3.3 artificial ear (blue) vs. a low noise Type 4.4 artificial ear (red) compared to the HATS self noise in the test room [3.3 artificial ear (dashed blue)/ Type 4.4 artificial ear (dashed red)].

special setting is named differently depending on the manufacturer; the realization strategies also differ. The typical setup for evaluating the talk-through performance of headsets is shown in Figure 3.

	ANC On	ANC Off	Transparency Mode
MOS-LE [9], [10]	1.9 MOS	2.4 MOS	3.0 MOS
Loudness Reduction (ISO 532-1) [11]	20.1 phon	15.7 phon	1.7 phon

Table 1: Listening effort measured with ABLE, according to ETSI TS 103 558: 07/2021 Methods for objective assessment of listening effort [9] for an intra-concha headset in different modes and related loudness reduction, measured with airplane noise.

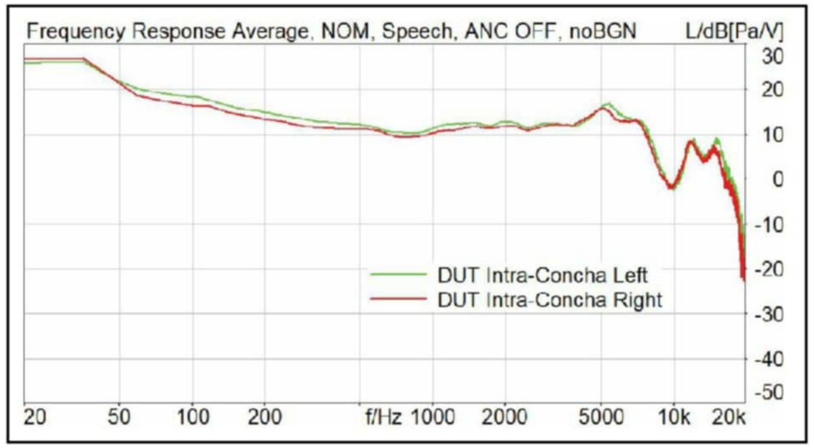

Figure 8: These are the response characteristics in Receiving for an intra concha headset.

HATS 2 is used to simulate the external talker in the noise field. A very suitable perceptual-based measure to determine the effort required to listen to a talker when wearing headsets is Assessment of Binaural Listening Effort (ABLE) [10] according to ETIS TS 103 558 [9]. On a scale between 5 (no effort required) to 1 (no meaning understood with any feasible effort), devices can be evaluated.

ABLE is a binaural method and determines the listening effort required based on the average user's perception. **Table 1** shows an example of an intra-concha headset with ANC on, ANC off, and transparency mode. With ANC "on," the listing effort is high as expected since outside noises, including talkers, are attenuated strongly. The associated loudness reduction of the outside noise is 20 phon. With ANC "off," the listening effort is less, but the loudness reduction is less as well. In "transparency mode," it can be seen clearly that the listing effort is relatively low; the measured MOS is 3.0. So, it is much easier for the user to listen to an external talker. However, it must be noted that the loudness reduction in this situation is just 1.7 phon. Consequently, the manufacturers of headsets typically leave it up to the users to choose the setting adequate for the situation. This combination of measurements provides the most relevant information for headset manufacturers to implement the talk-through functionality in their desired way.

References

[1] ITU-T Recommendation P.57: 06/2021, Artificial ears

[2] IEC 60318-4:2010, Electroacoustics – Simulators of human head and ear – Part 4: Occluded-ear simulator for the measurement of earphones coupled to the ear by means of ear inserts

[3] ISO 389-7:2019, Acoustics – Reference zero for the calibration of audiometric equipment – Part 7: Reference threshold of hearing under free-field and diffuse-field listening conditions

[4] ETSI TS 103 224: A sound field reproduction method for terminal testing including a background noise database

[5] ITU-T Recommendation P.58: 06/2021, Head and torso simulator for telephonometry

[6] HEAD acoustics datasheet: MDAQS – Multidimensional Audio Quality Score

[7] HEAD acoustics datasheet: 3QUEST 3- fold Quality Evaluation of Speech in Telecommunications

[8] ETSI TS 103 281: 05/2019, Speech quality in the presence of background noise: Objective test methods for super-wideband and fullband terminals

[9] ETSI TS 103 558: 07/2021, Methods for objective assessment of listening effort

[10] HEAD acoustics datasheet: ABLE – Assessment of Binaural Listening Effort

[11] ISO 532-1: 06/2017, Acoustics - Methods for calculating loudness - Part 1: Zwicker method

[12] HEAD acoustics datasheet: HMS II.3 LN HEC – Fullband artificial head with very low inherent noise and a human-like ear canal

Receiving Characteristics

For audio playback and speech communication, sound quality in receiving is vital. The basis for a physically correct measurement is a human ear reproduction and the accurate simulation of the average human ear impedance. The positioning of the headsets on the HATS has to be made as careful as for the ANC performance measurements. The basis for a perceptually correct measurement is the right choice of the HATS equalization in Receiving direction. The most often used equalization is the diffuse-field equalization (e.g., as defined in ITU-T Recommendation Head and torso simulator for telephonometry [5]).

When applied, the measured response characteristics should be almost flat—of course within certain limits and not considering the manufacturer's unique flavors in sound reproduction. A typical example of a frequency response characteristic measured in full-band mode is shown in **Figure 8**. Left and right ear characteristics match pretty well. A slight roll-off in the high-frequency domain, as well as a dip at 10kHz, can be observed. This measurement forms the basis for instrumental tests in receiving. However, no final judgment about speech or audio sound quality can be given based on this single measurement. Perceptually motivated evaluation procedures like MDAQS [6] for audio quality evaluation provide better insight into timbre, distortion, immersiveness, and overall quality. The results are presented on an mean opinion score (MOS)—a scale ranging from 1 (poor) to 5

Intra Concha Headsets	Timbre	Distortion	Immersiveness	Overall
H 1	2.5	2.9	2.6	2.3
H 2	4.7	3.4	4.3	4.4
H 3	4.4	3.4	3.4	4.1
H 4	4.7	3.3	3.9	4.3

Table 2: This is the Audio Quality (in MOS) measured with MDAQS [6] for different intra-concha headsets.

	Road	Café	Train
S-MOS	3.4	3.4	3.7
N-MOS	3.2	3.2	4.6
G-MOS	2.8	2.8	3.6

Table 3: Speech Quality (S-MOS), Noise intrusiveness (N-MOS) and Overall Quality (G-MOS) is measured with 3QUEST according to ETSI TS 103 281 [8] for an intra-concha headset, measured with different noises from ETSI TS 103 224 [4].

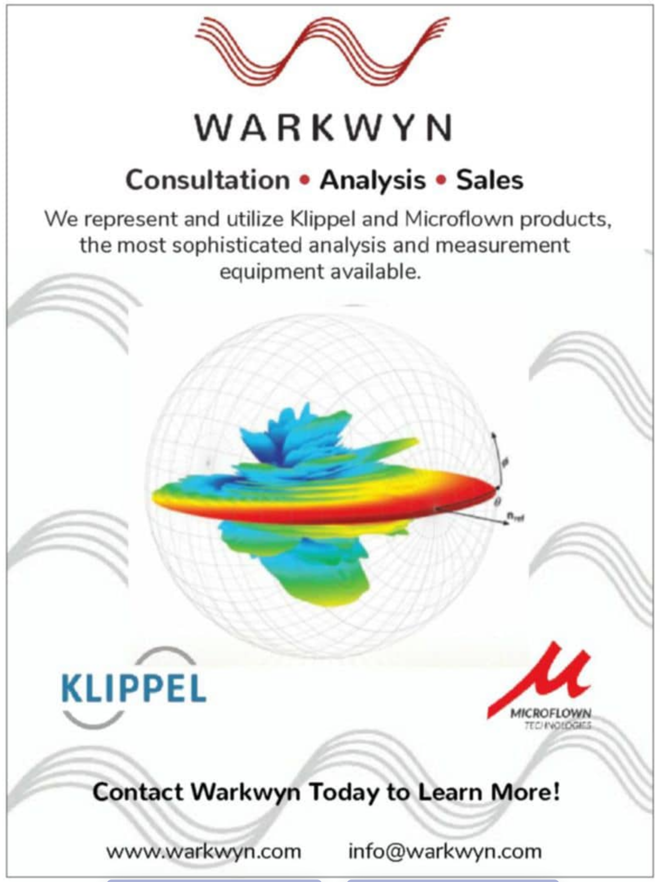

(excellent). An example for a comparison of four different intra-concha headsets is given in **Table 2**. Three of the headsets provide results higher than MOS 4 for the parameter timbre. The distortion quality is rated 3.3 to 3.4 for these headsets showing some weakness concerning this parameter for all headsets. Concerning immersiveness H2 shows clearly outstanding performance. H1 provides lower quality for all quality dimensions.

A measure providing perceptually well-correlating results for speech quality in noise is 3QUEST [7], [8], allowing us to determine the speech quality, the noise performance, and the overall quality separately. Again, the basis is using a diffuse-field equalized HATS equipped with a low-noise ear simulator replicating the human ear impedance correctly. An example result for one intra-concha headset with different background noises is given in **Table 3**. This headset works best with train noise while noise cancellation with road noise and cafe noise works less well. The speech quality for all background noises remains fairly high.

Conclusions

ANC headset testing requires a HATS with realistic artificial ears, low self-noise, and correct equalization. In addition, a setup for realistic background noise simulation and appropriate background noises are required. These measures can be combined with a second HATS, simulating an external talker. Various analyses are applicable for the different use cases of ANC headsets. Some examples of meaningful tests in receiving are shown in this article. Sending and conversational performance evaluation methods will be covered in a separate article. ax

About the Author

Hans W. Gierlich started his professional career in 1983 at the Institute for Communication Engineering at RWTH, Aachen. In February 1988, he received a Ph.D. in electrical engineering. In 1989, Gierlich joined HEAD acoustics GmbH in Aachen as vice president. Since 1999, he has been head of the HEAD acoustics Telecom Division and in 2014, he was appointed to the board of directors. Gierlich is mainly involved in acoustics, speech signal processing and its perceptual effects, QOS, and QOE topics, measurement technology, and speech transmission and audio quality. He is active in various standardization bodies such as ITU-T, 3GPP, GCF, IEEE, TIA, CTIA, DKE, and VDA. He is vice chair of the European Telecommunications Standards Institute (ETSI) Technical Committee for "Speech and Multimedia Transmission Quality" and was chair from 2016–2020. In 2021, he received an ETSI fellowship award.

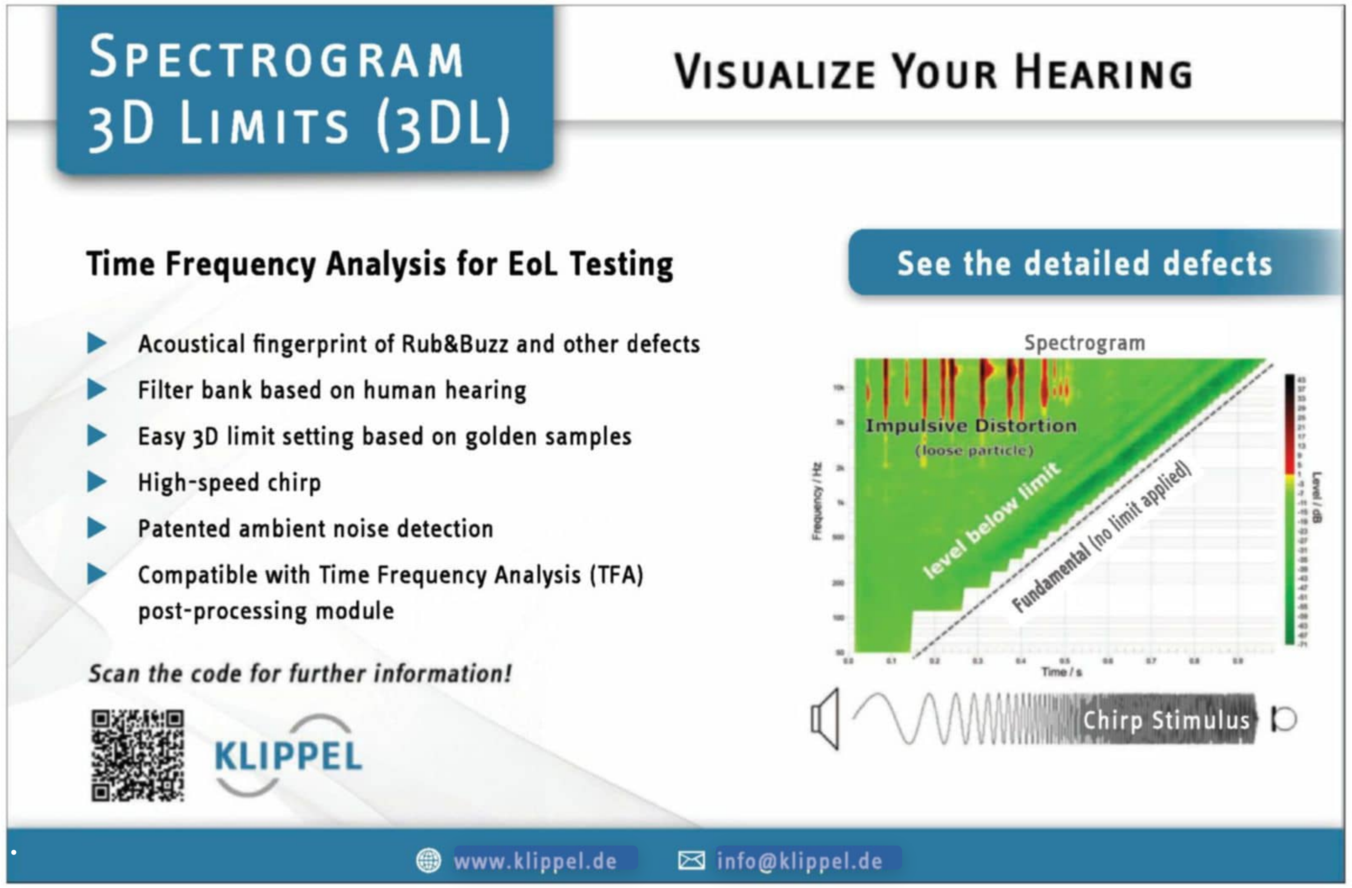

Building a Guitar-Controlled Synthesizer

LFO and ADSR

By
Ethan Winer

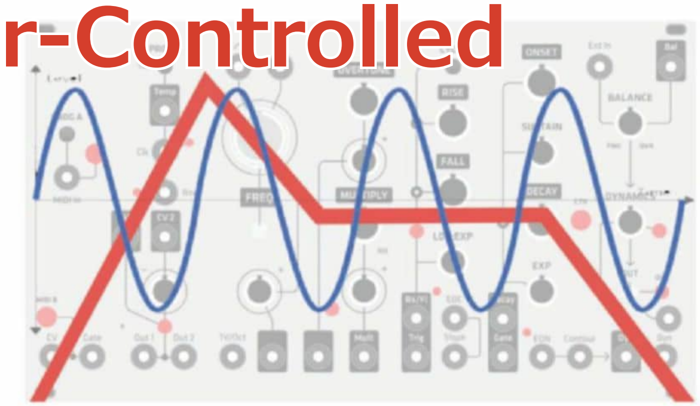

Ethan Winer's two previous articles described circuits he developed for his current project that uses an electric guitar instead of a keyboard to control an analog synthesizer. This article continues with two components that are essential for all analog synthesizers: The Low Frequency Oscillator (LFO) and the Attack Decay Sustain Release (ADSR) envelope generator. Both of these circuits generate voltages that control other modules in the synthesizer.

A Low Frequency Oscillator (LFO) is used to modulate the pitch of notes for vibrato and sweep effects, or vary the volume cyclically to create tremolo, or change the pulse width of a square wave to add animation, or sweep the frequency of a voltage controlled filter. Typical rates for an LFO range from slower than one cycle per second to around 10 cycles or even a little faster. Vibrato and tremolo effects are usually around three to four cycles per second, though sweeping an oscillator's frequency at a slower rate can sound like a police siren. I say cycles rather than Hertz because the audible effects of an LFO are cyclical. Some LFOs can output higher frequencies in the audible range above 20Hz to generate AM and FM sidebands for additional tone colors, or to serve as an extra audio oscillator.

An Attack Decay Sustain Release (ADSR) is an envelope generator, which is a fancy way of saying it adjusts the volume or other note qualities over time. Each time a note starts, the ADSR outputs a specific voltage sequence that modifies the parameters of the destination module. Both of these circuits control sound qualities, which can be difficult to convey in words. Under Resources is a link to my YouTube video "How to program analog synthesizers" with many audio examples, including how an LFO and ADSR can vary the sound over time.

The LFO

Most LFOs offer several waveforms including square, pulse, triangle, and sawtooth, as shown in **Figure 1**. A square wave can be changed to a pulse by varying its duty cycle, and a triangle wave can likewise become a sawtooth. A triangle wave is useful for vibrato because it raises and lowers the pitch at the same rate. Adjusting the output voltage of the LFO controls the strength of the effect. A sine wave is more difficult to generate, and when used as a modulator in an LFO it sounds about the same as a triangle wave. So there's little need to offer a sine wave.

The pulse and sawtooth waves can also be biased in the opposite direction, so the pulse is positive for longer than negative, or the sawtooth rises more slowly than it descends. As audio, such waveform pairs would sound the same because they have the same harmonic content. But as control voltages for pitch, the sawtooth pair would sound like two different sirens—slow up and fast down, or vice versa. Using a pulse for volume would make either the loud part or soft part longer or shorter. Using a square or pulse wave to shift the frequency will toggle between two note pitches (e.g., a European police siren) rather than glide smoothly from one to the other as happens with a triangle or sawtooth wave.

Earlier I mentioned creating AM and FM sidebands. When the volume or frequency of a musical note is modulated at an audible rate, additional frequencies are created related to the sum and the difference of the source and modulating frequencies. Depending on the frequency ratios, depth of modulation, and harmonic content of the source and modulating frequencies, this can create a variety of interesting timbres from warm tones to harsh metallic clanging effects. AM and FM are a staple of digital synthesizers, but they can be created just as easily using analog circuitry. For fun, the Project Files found in the Supplementary Materials section of the *audioXpress* website has links for two pairs of 2-second .wav files I created using both AM and FM to modulate a swept frequency at a fixed rate. One pair modulates a swept sine wave and the other sweeps a sawtooth wave. The source waves sweep from 400Hz to 600Hz while the modulator stays at a constant 80Hz.

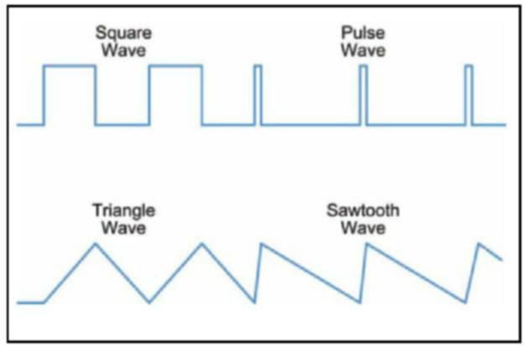

Figure 1: Most LFO circuits create these four basic waveforms. A pulse wave is based on a square wave, but with a duty cycle other than 50%, and likewise for the sawtooth derived from a triangle wave.

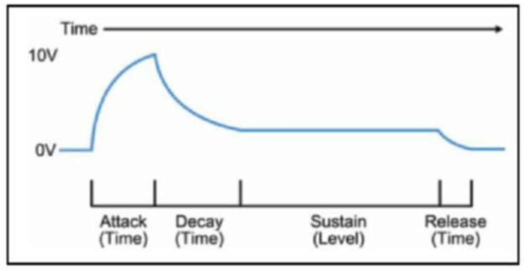

Figure 2: This shows the output voltage of a typical ADSR sequence every time a note is played.

The ADSR

ADSR—Attack, Decay, Sustain, and Release—are parameters set by the user via potentiometer knobs on the front panel. Some synthesizers have additional parameters, such as an initial delay before the ADSR sequence begins, or a second set of attack and decay times that happen after the first pair completes. But most ADSRs offer these four basic parameters. Note that the Attack, Decay, and Release settings control the length of each event, but the Sustain knob controls the sustained voltage level before the final Release.

Figure 2 shows how the output voltage changes over time for a typical ADSR setting, and **Figure 3** shows the result as applied to the volume of a 100Hz sine wave. This sequence begins whenever a note is played, whether on a keyboard or, in this case, an electric guitar. When a note starts the ADSR's output voltage begins to ramp up until it reaches 10V. Of course, 10V is arbitrary based on the synthesizer's design. Further, the amount an ADSR affects its destination module can be varied by its output amount, or by an input level control at the destination module.

Once the voltage reaches its maximum at a speed set by the Attack time, it begins to fall at a rate set by the Decay control. But rather than decay back to zero volts, it instead drops down to the Sustain level. When used to control volume, this lets you create a short note burst for greater articulation, then settle to a lower volume for the duration of the note to better blend with the rest of the music. This is a common playing technique for brass and other acoustic instruments. When the player stops the note, the volume falls back to zero at a rate determined by the Release setting.

LFO Circuit Details

As with my previous articles in this series, the circuits shown here were developed using the freeware LTspice virtual circuit simulator. This program from IC maker Analog Devices offers a fully stocked parts bin with a virtual oscilloscope, Bode and FFT frequency analysis, signal generators, and more. You draw the schematic, then click Run to watch it work. Once the circuit performs as expected you can buy the parts and build it for real. This is much faster and easier—and less expensive!—than using a breadboard, especially with a large project such as this synthesizer. And you'll never blow up an expensive component accidentally. Links to the program's download page and support forum are under Resources.

Before diving into the LFO circuit it will be useful to introduce the concept of hysteresis. This applies a controlled amount of positive feedback to preserve a system's state, and in practice means "Don't turn off the furnace until you get slightly above the temperature I set." Hysteresis resists change by requiring a condition to be exceeded, not just met, before the change finally occurs. Without hysteresis—typically a bi-metal coil spring attached to a switch and magnet—your furnace or air conditioner would cycle on and off every few seconds as the temperature varied ever so slightly above and below the thermostat setting.

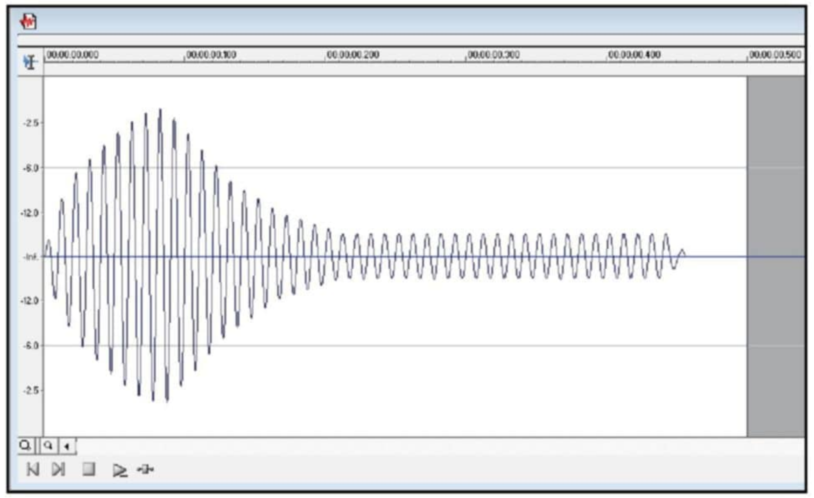

Figure 3: When the ADSR voltage sequence shown in Figure 2 controls the volume, the destination audio level changes in step with the voltage.

Figure 4 shows the schematic for the LFO circuit. Op-amp U1 is configured as an integrator that charges capacitor C1 up then down over time, and that provides the triangle wave output. The triangle voltage feeds comparator U2 whose output switches from fully positive to fully negative each time the triangle passes through zero volts. This serves as the LFO's square wave output. However, R1 and R2 implement hysteresis, so the triangle voltage must exceed zero volts in both directions before the comparator's output flips.

With nearly equal resistor values for R1 and R2, the triangle wave must ramp up almost to the +15V power supply before it overcomes the opposing -15V from the comparator's output. That finally forces its own plus input past zero volts (ground) at the minus input. When that happens the comparator's output flips polarity and capacitor C1 begins ramping in the opposite direction. Without hysteresis the triangle's output voltage would be extremely small and at a very high frequency due to the comparator constantly switching states around zero volts. You can see how this works by reducing R1 while monitoring the output and plus input of U1.

Note that LTspice doesn't have a potentiometer component. So in both the LFO and the ADSR circuits you'll see some potentiometers set up as resistor pairs, with separate values for each side of the pot's variable wiper. For example, R3 in Figure 4 is almost fully counter-clockwise to set a low frequency using 9k + 1k. If you play with the circuit in LTspice, change both resistors so they total the stated value.

Potentiometer R3 controls the amount of current available to charge C1, so that sets the LFO's frequency. There are also limit resistors R4 and R5 that prevent too much or too little current for the circuit to operate properly. Potentiometer R6 uses a pair of diodes to vary the capacitor's up and down ramp times independently. The waveform drawings to the left of R6 show how the duty cycle can be varied over a wide range in either direction. It's not practical for a single knob to control a frequency range much larger than 60-to-1 used here. So the final synthesizer will have a switch to select C1 as either 0.033µF as shown, or 0.0033µF to allow frequencies 10 times higher.

The last LFO item is the LTspice directive .ic, which sets an Initial Condition at any point in a circuit. Due to the vagaries of power supply turn-on voltage, LTspice can't always know what voltage will appear initially at every node in the circuit. Depending on the values of C1 and R3, without .ic this circuit won't start to oscillate until several seconds have passed. Forcing 10mV into U1's minus input begins ramping C1 to kick start the circuit. Note that the label In is not an input to the circuit, but rather the circuit node where LTspice is to insert the initial voltage. You don't need to name .ic destination nodes, but doing so is good practice because the internal node numbers can change as a circuit is edited. Naming the node guarantees that .ic will always be applied to the correct place.

ADSR Circuit Details

This ADSR was first used in the monster analog synth my friend Leo Taylor and I built in 1970,

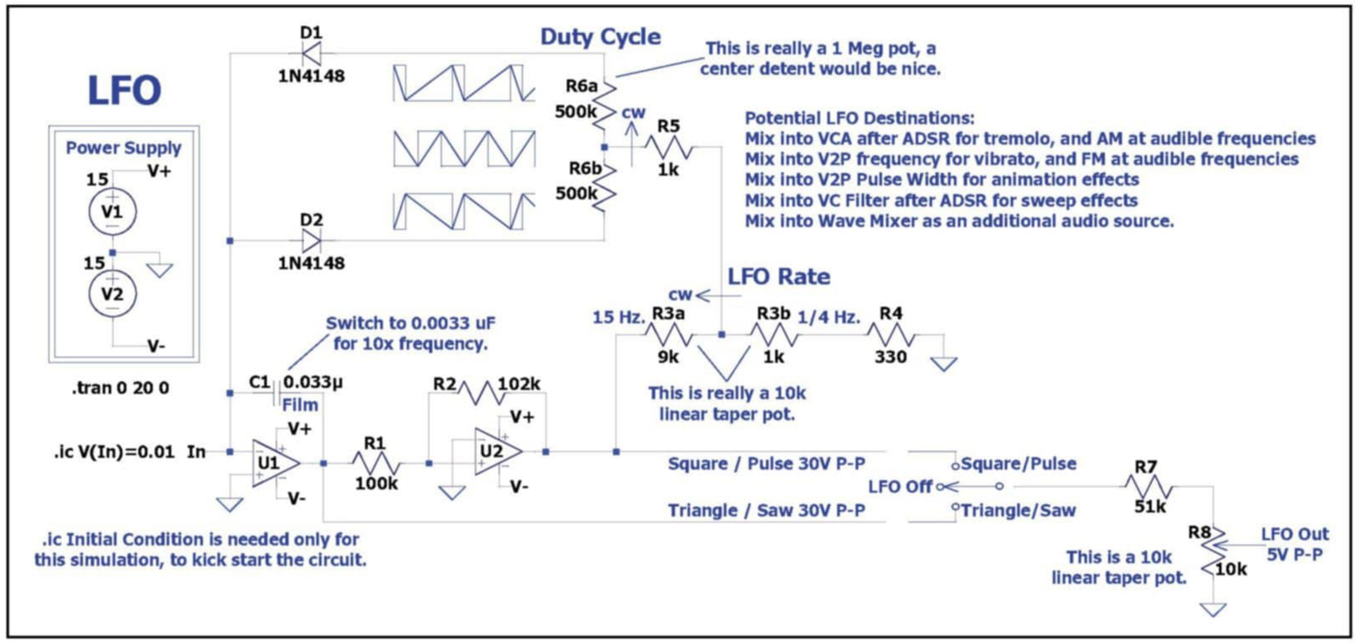

Figure 4: This LFO is based on an op-amp integrator having independent time constants for the positive and the negative directions.

About the Author

Ethan Winer has been an audio engineer and professional musician for more than 45 years. His Cello Rondo music video has received nearly 2 million views on YouTube and other websites, and his book *The Audio Expert* published by Focal Press, now in its second edition, is available at amazon.com and Ethan's own website ethanwiner.com. Ethan is also a principal at RealTraps, a US manufacturer of high-quality acoustic treatment.

mentioned in the April 2022 issue of *audioXpress*. If you have LTspice and run the ADSR.asc file from the Supplementary Materials, you can probe various places in the circuit to better follow the descriptions that follow.

In this guitar synth a Pick Detector (not yet shown) outputs a brief positive pulse every time a note is picked. So the first element in the ADSR is inverter A1 that reverses the polarity to negative as needed by the 555 timer's Trigger input to start the timing cycle. This trigger pulse also resets the timer, in case a new note is played before the Attack portion of the current note completes. Once triggered the timer starts charging capacitor C4 which, along with potentiometer R3, determines the Attack time. With the 80K and 2.2µF values shown, this is about 220mS. At the end of that period, the timer discharges capacitor C4 back to zero volts to begin another cycle when the next note is played. The remaining segments of the ADSR sequence also depend on C4 remaining at zero volts.

The input trigger also sends a brief pulse through transistors Q3 and Q4 to discharge capacitor C5 (**Figure 5**). This is the main capacitor that establishes the ADSR's behavior. A film type is used to increase linearity, and avoid voltage sag due to current leakage during the Sustain portion of the cycle. Although C4 sets the Attack time, C5 follows along because it's tied to the output side of diode D2, which pulls it up in sync with C4.

Note the second "digital" input at the lower left that stays high while the incoming note is active. That comes from a noise gate elsewhere in the synthesizer that monitors the incoming notes from the electric guitar. When a note fades completely, or is stopped by the player, the gate shuts down telling the ADSR to begin its final Release phase. Even though the incoming guitar note has ended, the synthesizer's oscillators continue allowing the note to fade out over the Release time.

Resuming the circuit description, the rising Attack voltage at C4 is buffered by op-amp U1 to

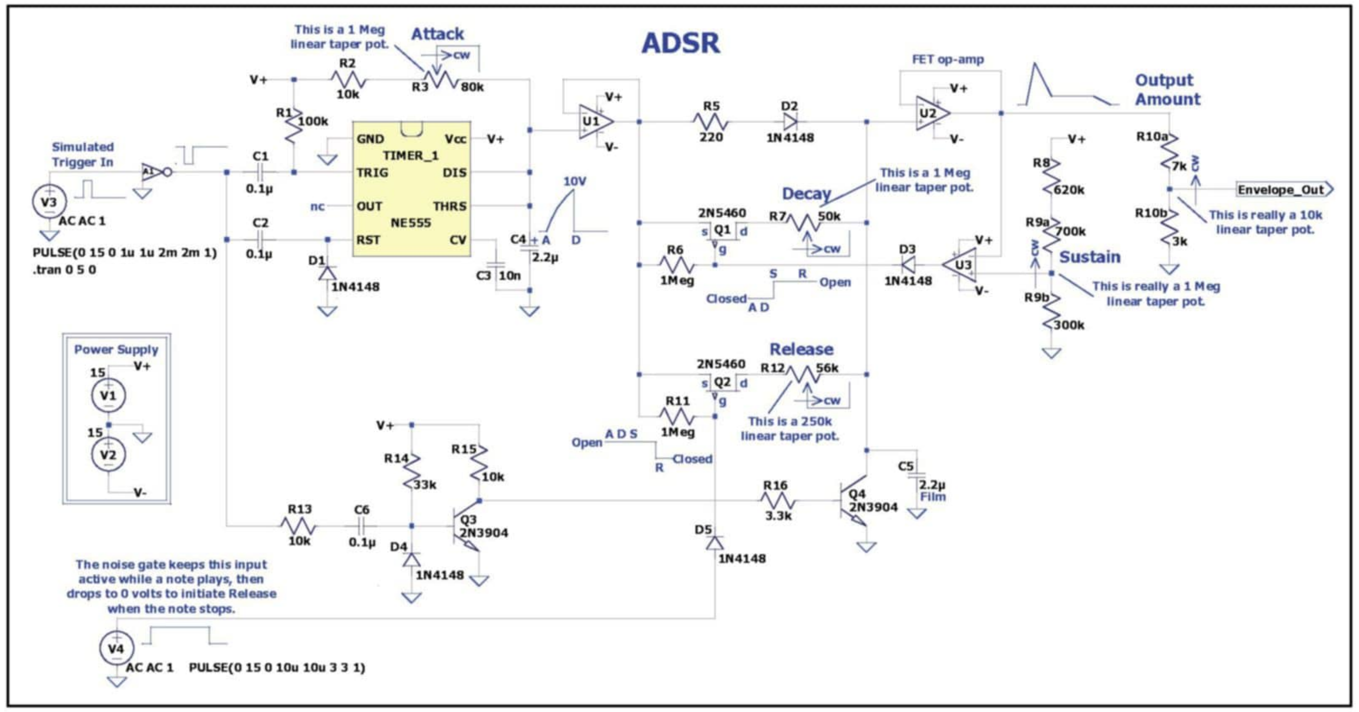

Figure 5: The ADSR uses FET switches to drain capacitor C5 independently for the Decay and Release segments of its cycle.

isolate C4 from the circuits that follow. That voltage is passed on to the output through another buffer, U2, but through diode D2. With the diode, the Attack cycle can only raise the output, but not drag it back down to zero volts when the timer finishes and resets C4. Instead, FETs Q1 and Q2 serve as switches that pull C5 back down to zero volts, first at the Decay rate, then later at the Release rate. In this case, zero volts is the output of U1 rather than an actual ground connection. An FET op-amp is used for U2 because its very small input bias current avoids draining C5 during the Sustain portion of the cycle.

When a FET is used as a switch, the voltage between its source and gate determines the switch state. The 1M resistors R6 and R13 tie each FET's gate to its source to hold the switches in the closed state (conducting) until the gate is forced positive relative to the source. P-channel junction FETs, such as the 2N5460 used here, have a resistance of about 300Ω when closed. Then when the gate is at least 3V to 4V more positive than the source the switch becomes an open circuit. Note that diodes in series with each FET gate let the comparator and noise gate pull the FET gates in the positive direction only, to force them open.

When the Attack portion completes, the output of U1 returns to zero, which begins to discharge C5 through FET switch Q1. This switch remains closed until C5 has drained down to the Sustain voltage. At that point comparator U3 flips positive and opens switch Q1, so the capacitor continues to hold its charge at that voltage. When the note finally ends, switch Q2 is closed by the noise gate, which discharges C5 the rest of the way down to zero. Again, the output of U1 is the zero volt discharge path instead of ground.

That's All, Folks

I hope readers enjoyed seeing these circuits as much as I enjoyed describing them. As "digital" and DSP have overtaken the world of audio, it's refreshing to know that there's still room for purely analog designs. Op-amps especially are amazingly powerful and versatile circuit building blocks, and they're still at the core of many audio products including microphone and phono preamps, active crossovers, and of course analog synthesizers!

Project Files

To download the links for two pairs of 2-second .wav files created using both AM and FM to modulate a swept frequency at a fixed rate, the LTspice circuit for the LFO, and the LTspice circuit for the ADSR, visit:
https://audioxpress.com/page/audioXpress-Supplementary-Material

Recources

LTspice,
www.analog.com/en/design-center/design-tools-and-calculators/ltspice-simulator.html

LTspice online support forum, https://groups.io/g/LTspice

E. Winer, "How to Program Analog Synthesizers,"
www.youtube.com/watch?v=NMF8F9z7Zr8

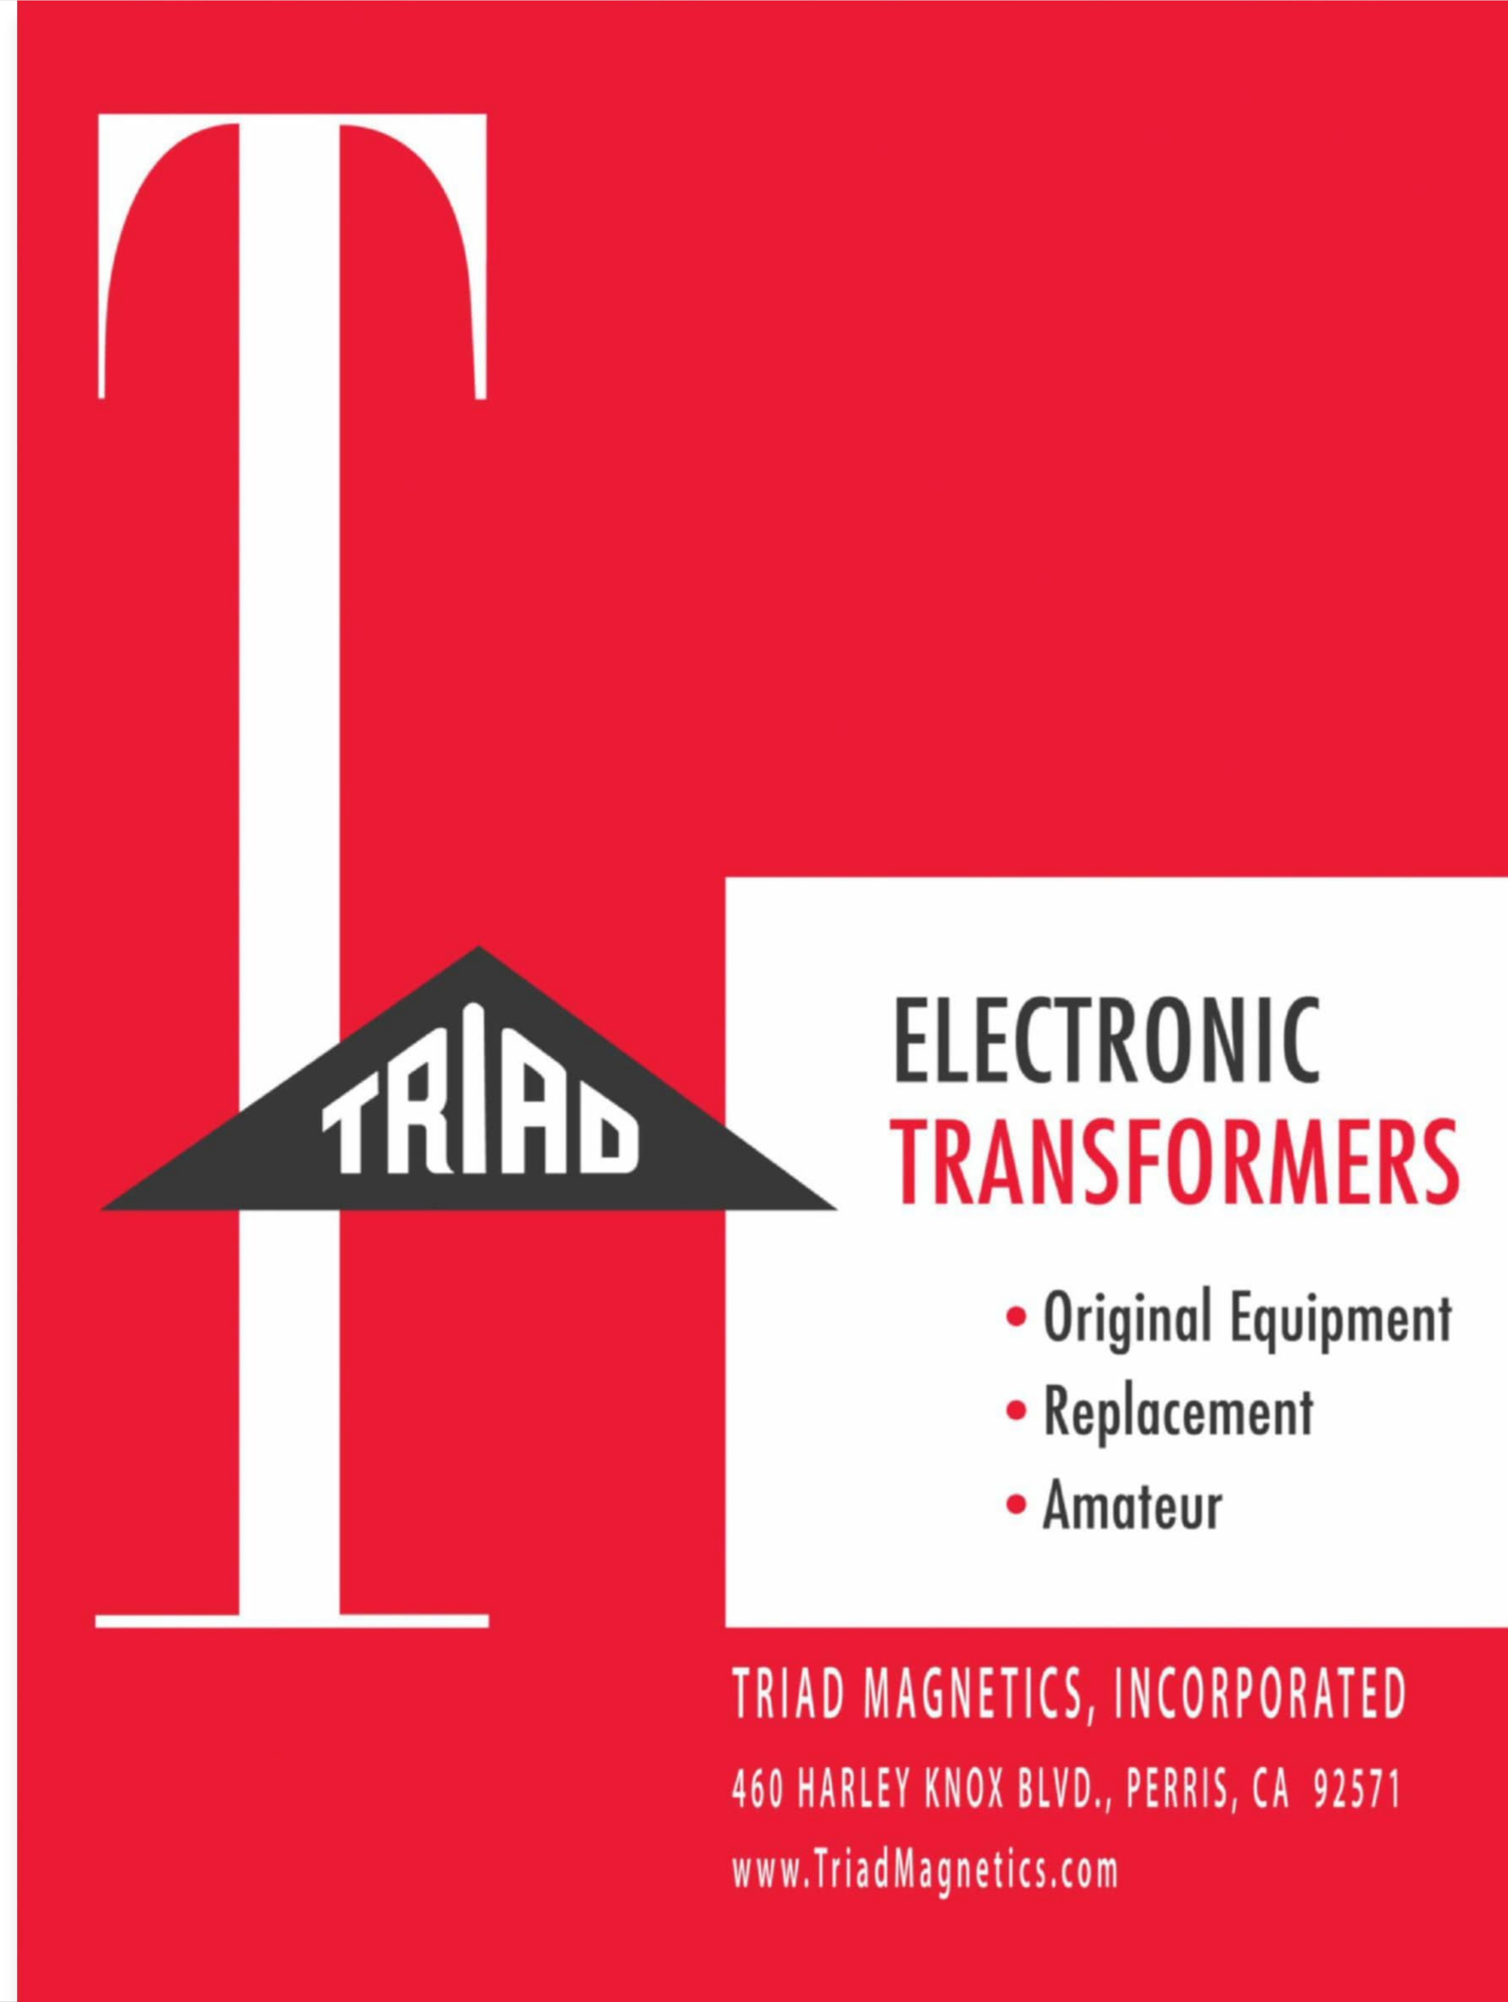

New KABD DSP Amplifier Boards from Dayton Audio

THE NEXT LEVEL OF DSP AMPLIFICATION

The Dayton Audio KABD line of amplifier boards give you unprecedented control for your next boombox project, active bookshelf speakers with DSP crossovers, 2.1 desktop system, or anything you can imagine. Featuring integrated digital signal processing, you can control your sound signature using free SigmaStudio software. With easy to configure solderless connections, drag and drop DSP, and predesigned examples, you are just clicks away from your own unique sound.

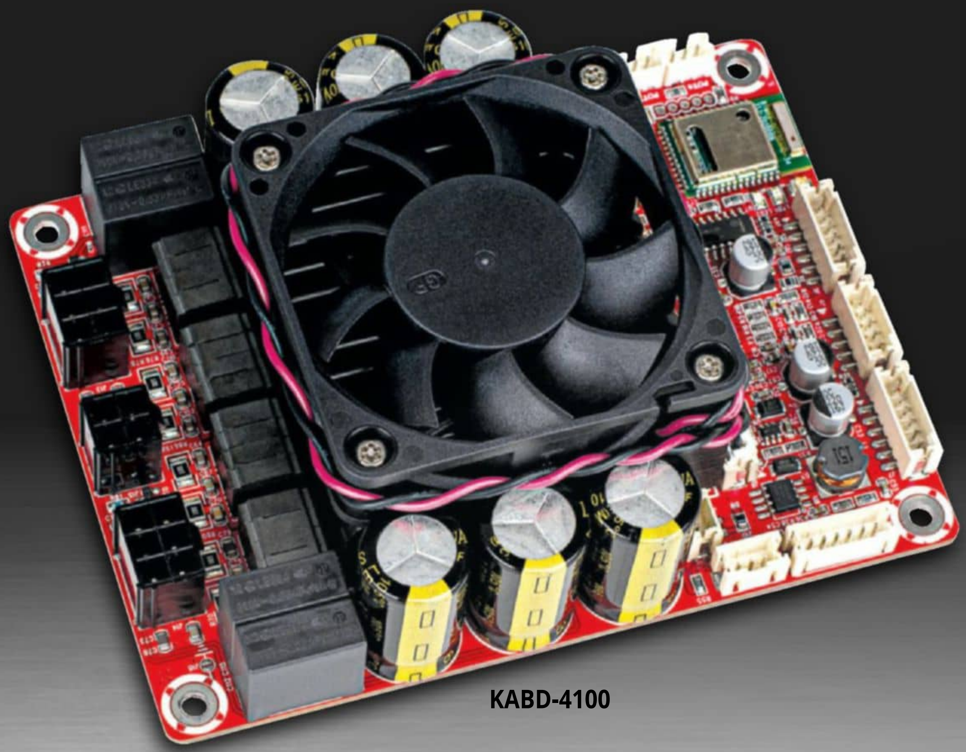

KABD-4100

What's New:

- All new versions: 2 x 30W, 2 x 50W, 4 x 30W, 4 x 100W
- Bluetooth 5.0 with aptX HD for high quality and long-range streaming
- Optimized standby mode and efficient amplification is perfect for battery-powered projects
- External potentiometer ports allow easy integration of volume, tone control, filters, and more for any project

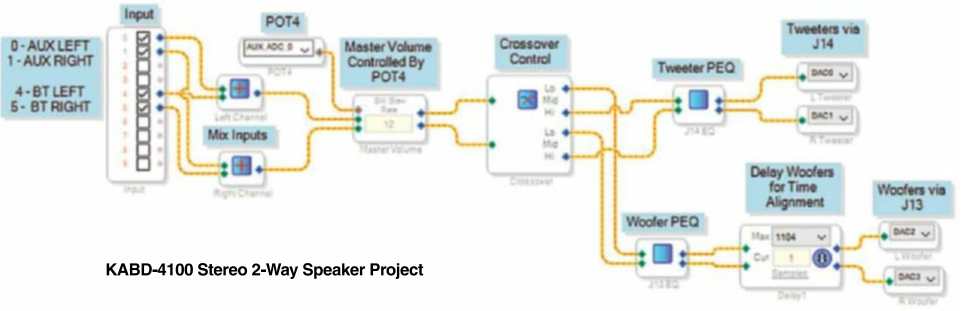

KABD-4100 Stereo 2-Way Speaker Project

DAYTONAUDIO
daytonaudio.com

Distributed by:
PARTS EXPRESS
AUDIO COMPONENTS & SOLUTIONS
parts-express.com 1-800-338-0531
Custom OEM inquiries welcome

Scan here to learn more about the Dayton Audio KABD Amp Boards